BITS

computing & communications news

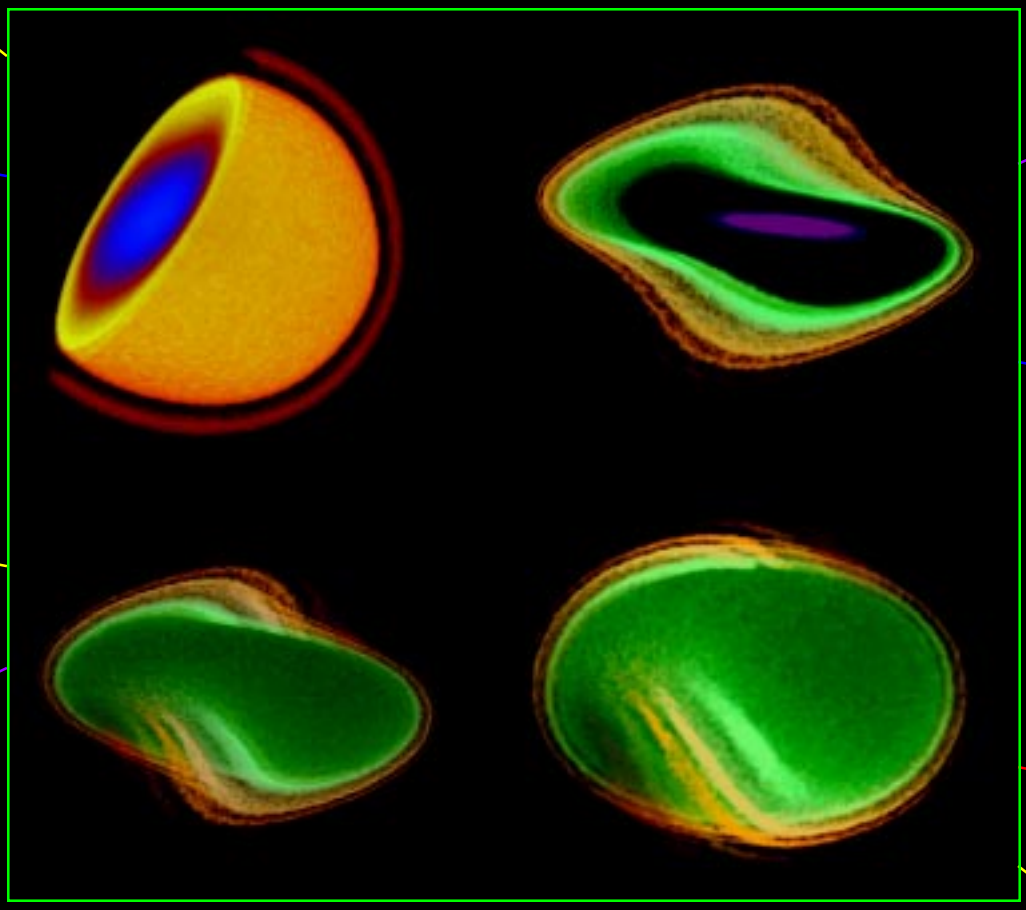

Inside

Infrastructure
Electronic Information Protection Regimes

Desktop
Software Project Management Series: Adjusting for Reality: Mitigating Uncertainty in Projects

Information Systems
Aha! An Interface to the Best of LANL Science and Technology Web Resources

High-Performance Computing
Tecolote: An Object-Oriented Framework for Hydrodynamics Physics

What's Happening
Telling Your Story in BITS: Sharing Your Expertise and Energy

Los Alamos
NATIONAL LABORATORY

August 1999

Front cover: **This computer-generated image shows four volume-rendered particle density fields based on a 300-million particle linear accelerator simulation. The output from such simulations contain three coordinates and three momenta for each particle at various locations along the accelerator. The upper left-hand image is based on the particle coordinates at one location. The remaining images show different views of a density field that includes a component of the particles' transverse momentum. The ability to visualize and analyze multiple representations of the six-dimensional output is essential to gaining insight from the enormous data sets produced by large-scale accelerator simulations. The data were produced using the code IMPACT (Integrated-Map and Particle Accelerator Tracking code), developed by a team involving LANSCE-Division, T-Division, the Advanced Computing Laboratory, and UCLA. Particle accelerators have led to fundamental discoveries in high energy and nuclear physics, materials science, and biological science. The design of the next generation of accelerators will require a new level of modeling as accelerators push the frontiers of beam intensity, beam energy, beam control, and system complexity. By making use of the latest high performance computing resources, and developing software and algorithms targeted to those resources, researchers are revolutionizing the design, construction, commissioning, and operation of next-generation accelerators. This science and visualization research is part of the Computational Accelerator Physics Grand Challenge project. For more information about the visualization science, contact Visualization Research Scientist Patrick M^{C}Cormick, Advanced Computing Laboratory, at (505) 665-0201, or e-mail pat@acl.lanl.gov. For more information about the Computational Accelerator Physics Grand Challenge, contact Principal Investigator Robert Ryne, LANSCE-1, at (505) 667-2839, or e-mail ryne@lanl.gov.**

Produced by the Computing, Information, and Communications (CIC) Division
Design: Donald Montoya and Julie Medina (CIC-1)
Managing Editor: Denise Sessions, denise@lanl.gov (CIC-1)
Illustration: Dave Delano
Photography: Gary Warren (CIC-9)
Printing: Imaging Services Group (CIC-9)

Attention Subscribers: There are two ways you can update your subscription data. E-mail bitsupdate@lanl.gov, or fill out the form inside this issue. Internal LANL subscribers please note that until we have a mechanism in place to verify our subscription database with the LANL Employee Information System we will rely on you to update your mailstop. Thanks for keeping us updated on subscription data.

Table of Contents

Information Systems **2**
Aha!–An Interface to the Best of LANL Science & Technology Web Resources 2
More than a New Look & Feel: The New LANL Badging System 5

Desktop Computing **7**
Adjusting for Reality: Mitigating Uncertainity in Projects 7
Your Data is Gone . . . Now What? 11

Infrastructure **13**
Electronic Information Protection Regimes 13
Web Content Accessibility: New W3C Guidelines Have Arrived 20

High Performance Computing **23**
Lecture Review: Next-Generation Chips, Processors, Transistors, and Wiring 23
Tecolote: An Object-Oriented Framework for Hydrodynamics Physics 26

What's Happening **28**
New Electronic Resources Available at the Research Library 28
Research Library Resources 29
Technical and Advanced Technical Computer Training Courses 30
Telling Your Story in BITS: Sharing Your Expertise and Energy 31
BITS Subscription/Comment Form 33
1999 12-Month Index 35

aha!—An Interface to the Best of LANL Science & Technology Web Resources

by the aha! Team: Mona Mosier (CIC-14), James Mottonen (CIC-5), Katherine Norskog (CIC-1), and Kathryn Varjabedian (CIC-14)

aha! is a search engine and directory developed for searching LANL science and technology resources. The charge for building a science and technology portal grew from the interaction between a group of people working on improved methods for searching the Laboratory's Web site and Deputy Director Bill Press. Press was concerned that people today who want to learn about science and technology at the Lab look to our Web site, but find it hard to get to the information.

Two Versions of aha!

This charge matched the desire to provide better searching mechanisms for both inside and outside LANL, and aha! was born. aha! is a Yahoo®-style directory and search engine. There are two versions of aha!—one for the public Web pages and one for LANL staff. The public version (developed first and previewed at LANL's IntraLab 99 conference) focused on resources accessible to the outside world. This version is incorporated into a new Science and Technology portal page off of the public home page. (See Fig.1.) In contrast, the internal LANL version includes additional resources of interest to LANL staff, such as ES&H, security, administrative and personnel information. aha! is not trying to categorize the entire LANL Web site; instead its goal is to categorize the best science and technology work being done at Los Alamos so that the public can readily locate information and resources most used by LANL staff.

The user can approach either version by performing a keyword search on the aha! database or by browsing through the category tree. (See Fig. 2.) This model was chosen to address searching needs. The present LANL search engine searches the full text of the entire Web site, making it difficult to retrieve the most relevant information among all of the press releases, News-bulletin articles, e-mail, etc. This is a common problem with general automated search engines. Often users are overwhelmed with large amounts of unfiltered information. Research studies have also shown that users don't know what to type in a search box, or they don't know how to use the format a search engine expects. Partial matches, misspelled words, and lack of understanding of relevancy ranking are problems. Homonyms and words used in unexpected ways cause problems; for example a search on "dinosaurs" retrieved an article referring to aging steel mills.

A Yahoo!-Style Directory

An alternative to the automated keyword search engine is the categorized directory, of which Yahoo! is the most well known. In this type of directory humans evaluate information, categorize resources, and assign descriptions. The categorized directory provides users more focused and higher quality information and the opportunity to browse.

aha! entries contain title, URL, brief description, and occasionally additional keywords. Up to three categories are assigned, depending on the topic and specificity of the item. Category trees can be browsed. Or, a search of the aha! database will search the titles, descriptions and keywords. A search can be performed on all of aha!, within one aha! category, or all of the LANL Web. Each aha! search displays results from the aha! database and also indicates the number of hits on the entire LANL web, so that if you don't find what you need in aha! you can simply follow the link to the broader results. aha! results are alphabetically arranged by title, with the categories assigned to each resource indicated, so that you can follow on to categories that look useful.

The search engine includes automatic right-hand truncation, and use of capital letters will result in a case-sensitive search which is useful for acronyms. Although the system currently does not support Boolean operations, this capability and other advanced searching features will be implemented in the future.

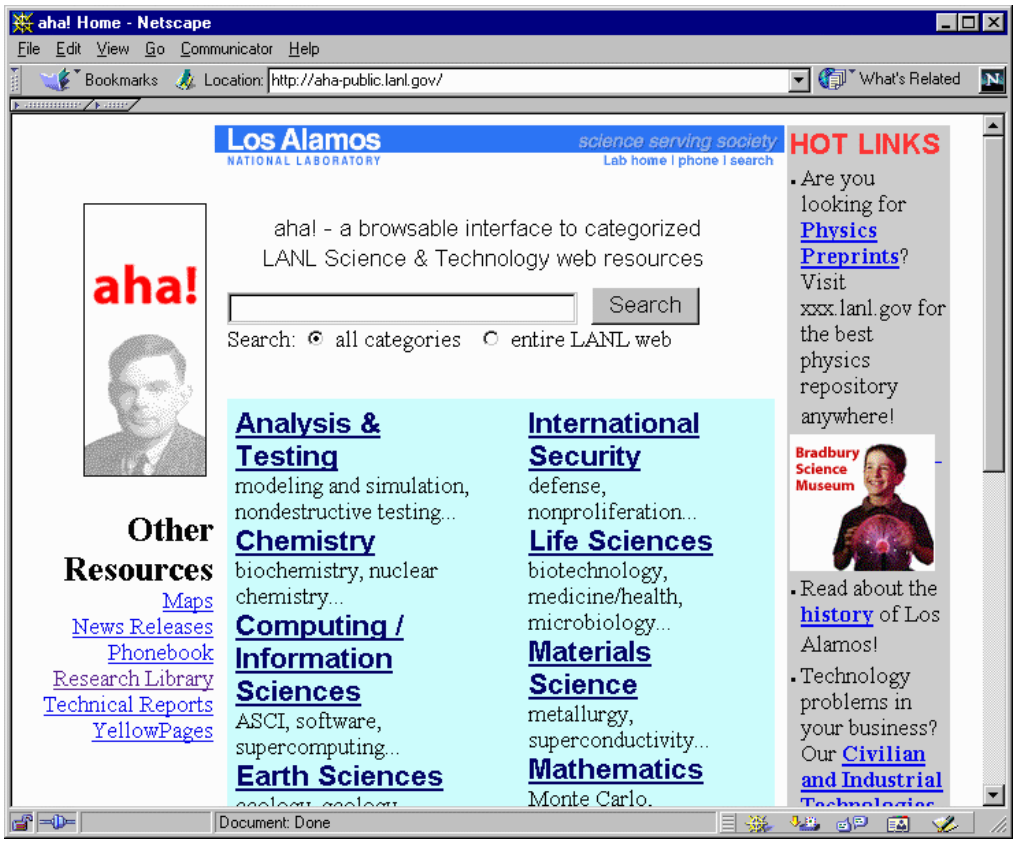

Fig. 1. The main page of the public version of aha!

Web Pages, Media Files, Hot Links and More

A variety of information resources are included in aha!: Web pages, technical reports, Los Alamos Science articles, Los Alamos patents, and media files. The Web pages represent and provide information on the science and technology work done at Los Alamos. The technical reports and patents are PDF files provided by the Research Library. The media files are RealMedia streaming video files which are selected from the LANL Media Server.

aha! also provides "Hot Links" on the main page and on each top-level category page. These links will be managed by editors within subject areas of the Lab and will give us a chance to make announcements related to a subject page or to call out links that we know to be HOT!

The category tree was developed by drawing on standard published thesauri and incorporating the ideas of subject experts throughout the Laboratory. It is structured to incorporate the majority of scientific and technical work being done at Los Alamos. This tree structure is open for suggestions–there is a "Suggest a Category" link off the aha! home page. Top-level tree headings are: analysis & testing, chemistry, computing/ information sciences, earth sciences, engineering, environmental sciences, industrial technologies, international security, life sciences, materials science, mathematics, nuclear sciences, organization/operations, and physics. These categories are further subdivided into more specific topics, and those topics will be subdivided again, as necessary.

Sites are gathered for inclusion in aha! in a variety of ways. The aha! team has searched the LANL Web for key sites. All Los Alamos patents are being entered, using information obtained from the Research Library's online catalog. Highly used technical reports and all new LA series reports are being added. Informal Los Alamos Unclassified Reports (LA-URs) are not regularly added to the database. Authors who want to ensure that their reports are made accessible on the Web through the aha! database should send them through the editing process and have them made into formal reports.

Web pages are also added by suggestion. An "Add a Site" option is available in aha! for you to suggest your Web pages if they are not already there. All suggestions will be reviewed by the aha! team for appropriateness. Pages with scientific content, more than just a sentence mentioning the type of work done, are desired. You may register personal pages with aha! if the content describes scientific research relating to the individual's role at the Laboratory. Appropriate use of the LANL Web guidelines apply. The editorial policy is available on the "Add a Site" page. After submitting a site, the aha! team will notify you via e-mail that your site has been entered or inform you that the site does not meet the guidelines. Also, as a bonus, the e-mail will include HTML-coded metadata that you can insert into your Web pages. This will improve retrieval of your pages by most search engines, in particular the "all of LANL Compass" search engine.

The Team Seeks Input

The aha! team is seeking volunteers to help with revisions to the subject category tree and assist with the placing of sites into that structure. If you are interested in working with aha! to provide subject expertise please send e-mail to aha@lanl.gov indicating your area of specialization.

aha! was developed by Mona Mosier, James Mottonen, Katherine Norskog, and Kathy Varjabedian and sponsored by the Director's Office and CIC Division. James Mottonen (CIC-5) developed the programs, Katherine Norskog (CIC-1) designed the interface and is coordinator of the project, Mona Mosier (CIC-14) and Kathy Varjabedian (CIC-14) designed the category tree and are responsible for populating the database.

aha! can be found on the Web at http://aha.lanl.gov or http://aha-public.lanl.gov. The product is continually evolving as new sites are added daily, new categories are created, and new enhancements are made. Please check it out and send feedback on the product to aha@lanl.gov.

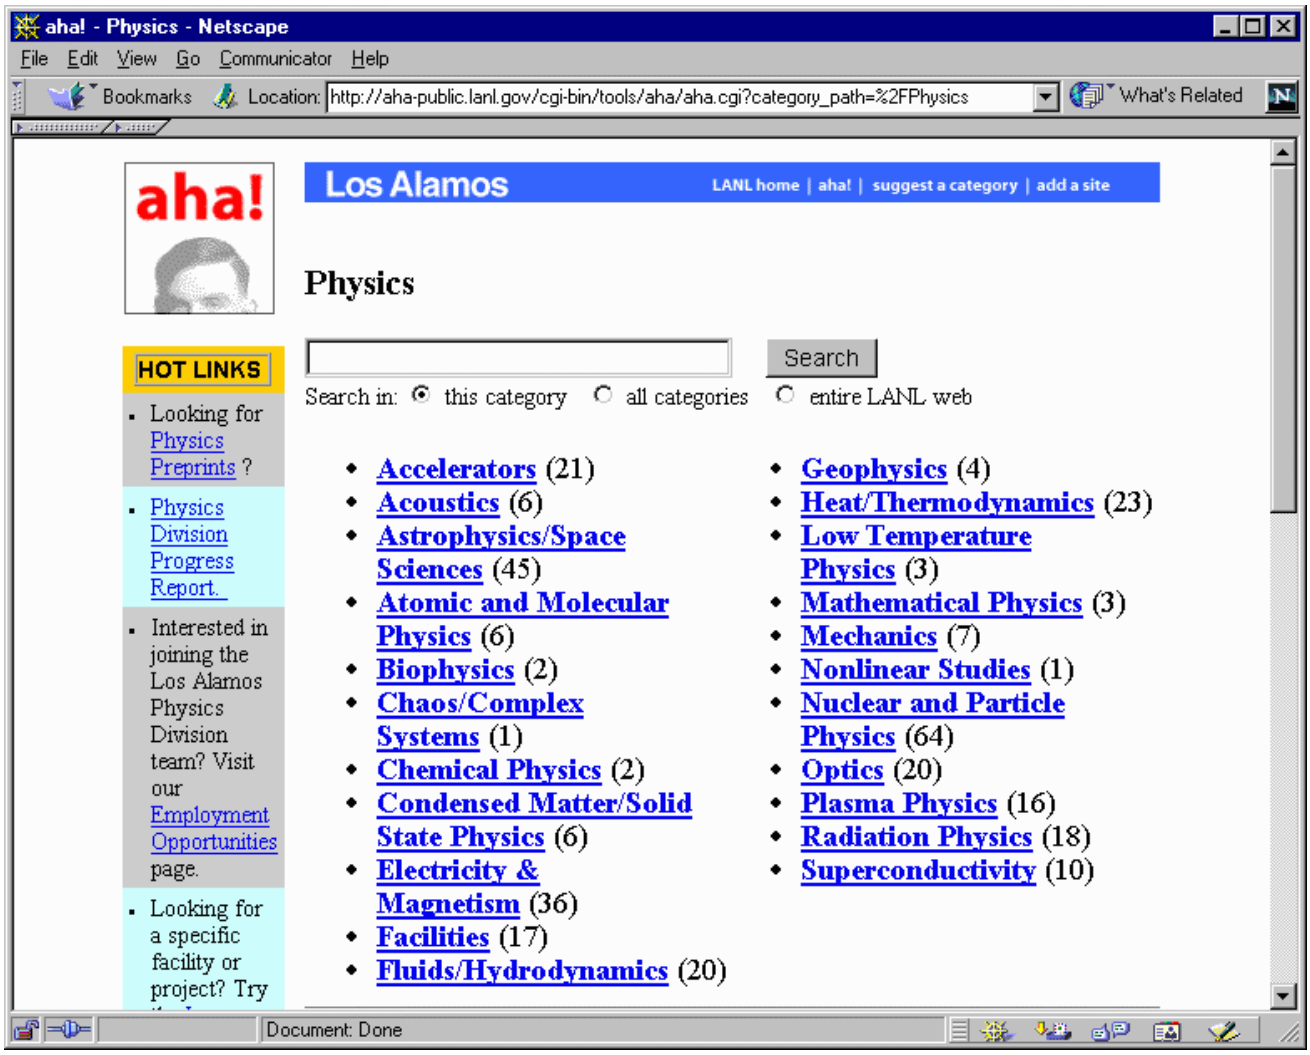

Fig. 2. A view of the physics link from the category tree.

More than a New Look & Feel: The New LANL Badging System

by Denise Sessions, BITS Managing Editor, CIC-6 Customer Service

What started out as a project to replace the current Laboratory badges with those made of more durable material has turned into a project that now reaches beyond the Badge Office and an outside vendor to several Laboratory divisions. Because the Badging Implementation Project is well worth examining from a reengineering point of view, BITS interviewed the project leaders: Ken Collins, Badge Office Team Leader (S-6), and Michael Smith, Business Information Systems (CIC-13) to find out all the details.

"The Laboratory has redesigned the entire system to link the badges and the badging process to databases in the Enterprise Information Applications (EIA, formerly known as Lab-wide Systems)," said Collins. The Laboratory seized the opportunity to add value to the project with these benefits. "In the new badging system, we've tried to address all of the issues that employees, Laboratory management, and DOE have raised in the past," Collins said. Instead of managing badge data on a stand-alone Badge Office system, the new project will reflect the use of Laboratory systems to access data such as badge status, terms of employment, clearance status, and security education.

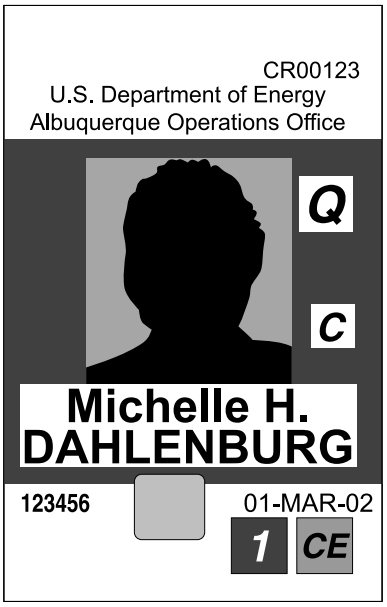

A Little History

Even though the Badging System Implementation Project is just beginning to loom on the horizon for most Laboratory employees, a couple of years of planning have been involved. Collins initiated the project in 1997 by soliciting for GIRE (**g**eneral **i**ndirect and **r**echarge **e**quipment) funding for a new badging system. After securing the funding, Collins convened stakeholders' meetings in the summer of 1998 to get the feedback he needed to design the system specifications. Stakeholder representatives from several divisions across the Lab participated as an implementation team. Led by Security (S) Division and Computing, Information, and Communications (CIC) Division, this team included individuals from other Laboratory divisions including human resources, nuclear materials technology, and nonproliferation and international security. Collins organized a team of key stakeholders to develop a procurement scope of work and select a vendor to create the badging system. In the midst of the procurement process, which used BUS-5's fairly new "Best Value" award approach, project co-leader, Michael Smith, came up with an idea that that changed the approach to this project. The approach had been to procure a system to produce new badges. His idea was that the new system should use the Laboratory's existing data in EIA. In this way a lot of data that already exists does not have to be re-entered. Using the Lab's corporate data repository will actually improve the Lab's data store by ensuring its consistency and conforming to Laboratory business rules.

Before anyone realized it, the project had taken on a greater scope. "Although there may be a few rough spots in the initial implementation phases," commented Smith, "a year from now, no one will imagine that we ever did things any other way."

Powerful Access-Control Tools

"By encoding employee information that connects the badges with the EIA system, badges will become powerful access-control tools," Collins said. Data that until now was available for badging only in stand-alone systems will be integrated in EIA. "What people will see is a badge that looks different on the front," Collins explained, "but the real power will be in the magnetic stripe on the back." Access to security areas and potentially, property-protected buildings and facility-specific access control systems also, will be

closely tied to data in the EIA system, including clearance status, accuracy of employee information and security education. For example, workers not current in their security education will be denied access to security areas. Similarly, access will be denied to subcontractors who change jobs but whose employers don't ask to have their clearances transferred or to workers who report a legal name change but don't change the names on their badges. "If a clearance is suspended, the badge readers in security areas throughout the Lab immediately will reject that badge," Collins said. Some entrances to security areas will continue to be staffed by guards, but within the next year badge readers will be installed at all entrances.

No one can argue that the new badge system will do anything short of improving security and general business practice at Los Alamos. Individuals and the people who support employees will be more accountable for fulfilling training requirements, meeting deadlines associated with security procedures such as completing clearance reinvestigation application forms, as well as for the timely entry of personnel-related data such as hire and pay status.

IA Principles in Action

The underlying power of the system is based on an integration of a number of existing and new databases within EIA. The system embodies the Information Architecture (IA) principles because it conforms to the client-server model. Although the "client" or, the customizable interface, is being written by Loronix Information Systems (the badging system vendor that won the competitive procurement) and some special printers are being procured, the Badge Office's "desktop" connection to the infrastructure conforms to IA principles. The principle is that data is entered one time—into the corporate repository—and from there the data is used multiple times. The Badge Office will use data from the corporate repository on a real-time basis. At the same time, other programs throughout the Lab use the same data. This process prevents errors in data entry and duplication of effort by reducing the amount of times data is entered and reentered into standalone systems.

Other major benefits are that badging and other security data become available Lab-wide, to those with the appropriate authority, by means of Data Warehouse reporting. Furthermore, the portrait images used on badges will be in EIA–again, with the appropriate authority, these images can be used for other work-related purposes (such as producing group organizational charts that include pictures, for sending out Access Denial orders).

By contributing the ideas to follow IA standards and use the data residing in the corporate repository, CIC-13 has been a key player. As the project co-leader, Smith explained, "This project is a way to leverage the power of integrated information systems. It is a very complicated project and it uses information infrastructure at the Lab in an unprecedented way."

Improvement of Security Procedures

This project contains very proactive security measures. It creates consequences for noncompliance with security procedures, thereby bolstering the Lab's ability to demonstrate compliance with security regulations and support "best business practices." With the use of real-time data, the Badge Office will be able to provide faster and more efficient customer service in the areas for which they are responsible. To help smooth the transition, the Badge Office has made a great deal of information about the Badge Implementation Project available to badge-seekers in many media forms, such as printed brochures and a Web site (badge.lanl.gov) or directly linked from the Laboratory's (internal) home page.

Future Aspirations

Despite all of the advantages of this new system, ultimately the Badge Office and its CIC-13 partner would like to get direct access to DOE clearance systems so the Lab doesn't have to continue to maintain their own clearance system. The value would be to validate data dynamically with improved integrity. What does it take to get people to grant access to their data? Perhaps as more and more people experience the consistent integrity of available data, they've accepted its reliability just as they trust the telephone system, for example. Smith commented that "people have been driven to the point of trusting the storage and integrity of computer networks." The Badging Implementation Project is an example of the Laboratory taking the opportunity to take advantage of their corporate database, and its networks and data management.

Bottom Line Impact

Oh, and by the way, the Badge Office expects the time to obtain a new badge will be reduced to a couple of minutes, in most cases.

For more information, see the Badge Office Web site: badge.lanl.gov

Adjusting for Reality: Mitigating Uncertainty in Projects

by Don Willerton, Group Leader
CIC-6, Customer Service

Software Project Management Series

This article is the second in a three-part series in BITS exploring how you can keep a better handle on your software development project to ensure it doesn't spin out of control or face cancellation before the product's release. One major reason for having difficulties in managing software projects is the intractable presence of uncertainty in developing software, causing traditional methods of project management to be ineffective. Recent developments in software project management theories and practices are addressing this uncertainty.

In the first article of this series, I identified internal management within software projects as a key to the success or failure of developing useable, effective, and used software. Central to this failure, the level of uncertainty in these kinds of projects defies using the more common assumptions, tools, and methods of traditional project management.

But, if we can't use traditional project management, what can we do?

It's a question that has many current software thinkers, researchers, and practitioners proposing ideas, methods, and practices that range from "just make a few modifications" to a wholesale "you have to change the very nature of the business!"

Here are three examples (among many) where really good, experienced, and successful software developers are proposing changing the way software project managers manage projects. Among the differences to watch for, notice how each handles the "unknownness" of the project, whether it's as tangible as project resource allocation or as intangible as knowing how the "real product" of the project should be defined.

Steve McConnell: Surviving Projects

Steve McConnell is a familiar face to some software teams at the Lab. Most recently, he's been teaming with two other software professionals to present a software development course and do some consulting with individual code teams. Steve is a remarkably skilled and experienced software professional, as well as an accomplished author. *The Project Survival Guide* is his third book.[1]

In giving context to McConnell's guidelines, consider "making things" from a process point of view. If traditional project management is concerned with making "something", then the underlying principle is that you can get better at it if you improve all of the processes used to make it. The basic idea is that you figure out how to make "something" by very formal, measurable, and recordable processes, analyze what went right or wrong when you're finished, improve the processes, and then do it again. After making the next "something", improve the processes again, etcetera, until you've optimized the processes to produce the "something" in a very reliable, repeatable, estimatable, and controllable way. This is why early bridges failed and later bridges did not. We learned something in between and built the new knowledge into the processes.

The Software Engineering Institute's (SEI's) Capability Maturity Model (CMM)[2] is the analogous approach to software projects. It is a process improvement-based model for producing software "somethings".

From this perspective, the project management techniques recommended by McConnell are not radically different. He's very serious about traditionally recognized software development practices such as technical reviews, staged development, change control, planning and plans, risk management, measured quality, predictive schedules, and estimated costs. Where he does differ from the CMM, which wants to drive out "mistakes" by process improvement, is the recognition of the inability of the project team to know everything about a project before it begins. He describes this inability by using the term "cone of uncertainty" in many aspects of the project. In estimating a project, for example, the effect is definite:

> "The cone of uncertainty has strong implications for software projects estimation. It implies that it is not only difficult to estimate a project accurately in the early stages, it is theoretically impossible." (Ref. 1, p. 32)

Because of that uncertainty, process improvement will never be able to proactively, or iteratively, eliminate mistakes. This means that

> "Success in software development depends on making a carefully planned series of small mistakes in order to avoid making unplanned large mistakes." (Ref. 1, p. 36)

With the assumption that you can't prevent mistakes, his recommended practices put a lot of effort into the beginning of the project.

- Have high visibility—a constant involvement of the customer, stakeholders, and management; everybody knows everything that's going on, all of the time.
- Relentlessly hold design reviews, quality reviews, code reviews, etc.—find the mistakes when they're cheap, instead of later on, when they can cost from 50 to 200 times the earlier cost.
- Give something to users in staged development—consistently and frequently release what you have so the user can be involved as soon as possible in seeing the product. "Staged releases force development teams to "converge" the software—in other words bring it to a releasable state—multiple times over the course of a project. Converging the software reduces the risks of low quality, lack of status visibility, and schedule overruns, . . ." (Ref. 1, p. 238)
- Work with 80% rather than 100%—"Try to complete 80% of the requirements before beginning architecture and 80% of the architecture before beginning detailed design. Eighty percent is not a magic number, but it is a good rule of thumb: it allows the team to do most of the requirements work at the beginning of the project while implicitly acknowledging that it is not possible to do all of it." (Ref. 1, p. 58)

With no apologies for a disciplined and structured approach, and yet appreciating the fact that uncertainty prevents you from doing all of the planning that you might like, Steve believes very strongly that

> "A successful project should be one that meets its cost, schedule, and quality goals within engineering tolerances and without padding its schedule or budget. After detailed plans have been made, the current state of the art supports meeting project goals within plus or minus ten percent or better. This level of performance is currently within the reach of the average software project manager." (Ref. 1, p. 4)

THE Airlie Council: Best Practices

It was in the late 1980s that "process improvement" began to reach a fervor. The SEI CMM, for example, was announced in 1987. Five years later, several program managers of the U.S. Department of the Navy had grown disillusioned with waiting on the promises of organizational "process improvement" in the area of software development. They formed the Software Program Managers Network (SPMN) to help develop more practical, applicable, and timely methods of project management.

In 1994, SPMN sponsored a Best Practices Initiative that initiated a three-prong approach to developing what they needed. The first effort, the AIRLIE Council, was a group of a dozen or so nationally recognized experts in the area of software engineering and management. The Council included such dignitaries as Roger Pressman, Ed Yourdan, Howard Rubin, and Tom DeMarco. The product of the Council was a list of "Nine Best Software Practices."[3]

The second effort was a number of focus groups that made further recommendations of "43 Best Supporting Software Practices" that connected to and supported the first list. The third effort was the formation of an oversight committee to watch the other two councils.

The movement to "best practices" was a specific and major departure from the "process improvement" orientation of both the traditional project management methodology (via the Project Management Institute[4]) and the CMM (via the SEI). It was a recognition that the uniqueness of software projects, the inherent uncertainty in both the product and process of management, and the volatile nature of project change, resisted the stability and repeatability needed for process-control-based management.

Jim Highsmith: Adaptive Complex Systems

If the percent of uncertainty in a project was on a scale with traditional project management on the left, Steve McConnell is close to the left, the Airlie Council is close to center, and project management based on adaptive complex systems is on the far extreme right of the scale. The newest and hottest method for software development comes by way of Jim Highsmith, a consultant and author, and is a by-product of some of the work done at the Santa Fe Institute. It takes the accounting of project uncertainty to a new level.

> "Why would we think high-change, high-speed, and high-uncertainty projects should also be predictable and controllable? They are boundable and manageable, but not predictable and controllable."[5]

In this case, then, the idea of "process improvement" is inappropriate: no projects will be substantially the same. A whole new mindset is needed, and it is centered on the need to develop a "learning project environment".

Jim's project would look like this:

- Project planning is very important. Develop overall plans, but don't drive for a low level of specificity. Instead, leave the specific planning for first "timebox" of the project.
- A "timebox" is a definite length of time, chosen by the project team and usually 3 to 4 months, where a "mini-project" of sorts is constructed. Project deliverables are set for the end of the timebox (if not a fully functional code, then a well-defined piece of it), metrics identified and measured, estimates recorded, and anything else that is determined to be important.
- The project is started with a fully dedicated, small, senior team.
- Towards the end of the timebox, all work is wrapped up for delivery, including metrics, reports, documentation, a working version of the software, etc.
- A full "post-mortem" exercise is held with all team members, customers, users, stakeholders, and managers. All parts of the project are reviewed and discussed.
- The major question to be discussed and answered is "WHAT DO WE KNOW NOW THAT WE DIDN'T KNOW THEN?"
- With that information, literally anything about the project can change. This might even include changing the very purpose of the project! (A traditional project manager would have a heart attack.)
- All estimates, metrics, directions, activities, milestones, and priorities are now reset for the next timebox and the project continues.
- Iteratively, openly, deliberately the project cascades through a series of timeboxes until the project deliverables are reached.
- It is usually true that the product of the project will be anything *except* what was initially expected!

A few of the assumptions behind this methodology include:

- Every project is a learning situation. Don't focus on individual performances, but how the team functioned and the quality of "learning" that took place.
- Every project is unique. Instead of going for "process improvement", base your progress on "people improvement". A major project goal will be to develop agile, adaptive, and emotionally intelligent teams that can flourish in high-change, high-speed, and high-uncertainty projects. "Learning about oneself—whether personally, at a project team level, or at an organizational level—is key to agility and the ability to adapt to changing conditions."
- Every project MUST continually include the customer, user, stakeholder, manager, or anyone else that has vested interest. They must "learn" just like the team, they must suffer through the hard decision-making, and they must adapt their expectations as the project unfolds.
- Every project depends on good postmortem reviews. "Postmortems tell more about an organization's ability to learn than nearly any other practice. Good postmortems force us to learn about ourselves and how we work." (Ref. 5)

This is entirely different from traditional process-oriented techniques. Specifically, Highsmith says:

> "Speed, change, and uncertainty don't succumb to optimizing practices. It goes beyond a shift from a linear to an iterative lifecycle—it goes to the heart of the difference between a fundamental belief in *adaptation* rather than *optimization*." (Ref. 5)

Summary

Some software project teams know exactly what should be done and how it should be done and when it should be done.

Some software project teams don't know, and can't know, any of the what, the how, and the when. Those projects may rely on information that is invented or discovered during the actual activities of the project.

Most software project teams are probably somewhere in between.

In all cases, trying to manage the projects without recognizing the existence and level of uncertainty, however much it might be, will be a prescription for a doomed project. Fortunately, some people are getting smarter and are learning how to mitigate the circumstances so that having successful projects can come closer to being routine.

In the next issue of BITS, in "GrassRoots Software Management: Simple Things To Do," I'll have a list of practices from a review of recent project management methodologies and best practices.

References

1. Steve McConnell, *The Project Survival Guide* (Microsoft Press, Redmond, WA, 1998).
2. Carnegie Mellon University, Software Engineering Institute, "Software, Systems Engineering, and Product Development CMMs®." Available: http://www.sei.cmu.edu/cmm/cmms/cmms.html .
3. Software Program Managers Network, Airlie Software Council, "Nine Principal Best Practices." Available: http://spmn.com/best_practices.html .
4. For more information on project management, see the Project Management Institute Web site. Available: http://www.pmi.org/ .
5. Jim Highsmith, "Managing Complexity," in "Application Development Strategies Newsletter, August 1998," Cutter Information Corp.

Your Data is Gone…Now What?

By Nikki Gaedecke, Information Architecture Project, Desktop Team, CIC Division

First in a Three-Part Series

Many people who lose data can revert to backup files. For individuals who fail to back up important documents, they risk being forced to recreate them should an "accident" occur. Imagine coming into work one morning to find everything on your computer is gone: no applications, no files, no Word documents, and no projects that you've been working on for the last couple of months. Gone, and you don't have your data backed up!

A cardinal rule in the world of computers is to back up your critical data: to copy files to a second medium such as a disk or tape as a precaution in case the first medium, your primary hard drive, fails. Although backing up data is so important, computer help desk personnel report that many Laboratory users do not back up their data. If they do back up their data, some are doing it incorrectly by backing up the wrong information or not understanding the backup tool.

Some reasons why people fail to back up data are because they:

- don't understand the importance of backup,
- don't know how to use a backup tool,
- forget to do backups because they don't have a routine for doing so,
- don't want to be bothered by what they perceive is a chore,
- never previously lost data, and
- don't know how data can be lost or corrupted.

Not only can the pain of recovering from data loss be high, but also the cost is high in time required to recreate the lost data.

Reasons to Back Up

- To fulfill computer security responsibilities Laboratory employees must protect their "information by offloading it to removable media and storing it in a locked office, desk, file, etc." *(Computer Security Profile Protections: http://int.lanl.gov/projects/ocsr/UserReg.html*).
- The Information Architecture (IA) Project and the Year 2000 Council strongly recommend that users back up data in anticipation of potential Year 2000 problems.
- To insure against the loss of important and, perhaps irreplaceable, data. For example, your electronic data could be as important as the promissory notes or birth certificates you might choose to put in a safe deposit box.
- As a matter of economical sense, it pays to avoid wasting time recreating lost data.

How Data is Lost

Data can be lost or corrupted for many reasons including the following:

- Software viruses can corrupt data
- Hardware failures can destroy data
  - — Hard disk failure
  - — Overheating
  - — Motherboard failure
- User error
  - — Choosing to use the File, Save option versus the File, Save As option
  - — Reformat floppy disk that contains data
  - — Accidental deletion
- Miscellaneous
  - — Software upgrades can result in data loss due to user error, installer problems, or conflicts with other applications
  - — Theft
  - — Disks or hard drive exposed to magnets
  - — Power outages and electrical spikes

What Should and Should Not be Backed Up

You should back up important files, documents, or e-mail messages, but not software applications. Software applications can be reinstalled, but your data is completely lost without a backup. For example, you would want to back up all your Microsoft (MS) Word documents, but do not back up the MS Word software application. Not only is backing up software applications unnecessary, it requires large amounts of disk space. For example, MS Word 97 is 46 MB, while a two-page MS Word document could be 37 KB. The average 1.44-MB floppy disk can hold several MS Word documents. On the other hand, you would need almost 32 floppy disks to hold the MS Word application itself.

Available Backup Tools

Factors, such as cost, convenience/simplicity, reliability, or automation, need to be weighed to determine which backup medium is right for you. Regardless of which tool you select, you must routinely back up your data and verify that you can restore the backed-up data. Not all individuals who have backup tools are actually backing up their data correctly.

Several backup tools are available at the Laboratory:

- ADSM (Adstar Distributed Storage Manager)
- removable media (removable storage drives)
- floppy disks
- tape drives

ADSTAR Distributed Storage Manager

ADSM is an IBM client/server software product that provides full and incremental backup to the customer. In other words, ADSM stores your backups on centrally managed servers so that you don't have to keep tape or diskette tapes of your own. On a daily basis, ADSM automatically copies a current image of your entire computer system.

CIC-7 offers this backup service to employees for a small monthly fee.

According to Michael Carter, NIS-3 system administrator, NIS-3 employees depend on the system administrators to make ADSM operate smoothly. NIS-3 is proactive about backups where the system administrator plays the key role: all Unix systems are backed up nightly via ADSM and Windows, and Macintosh systems are backed up via ADSM if the user requests it. NIS-3 system administrators will install ADSM and register customers. Based on the NIS-3 experience, Michael Carter believes that "*ADSM should not be recommended to the average computer user. Instead, ADSM works well for groups and individuals that have system administrators who can deploy it (July 8, 1999).*"

Removable Media

For Laboratory employees who do not have a system administrator, removable media such as Zip or Jaz drives, is another alternative. Removable media operates just like your regular floppy disk drive (a:\) except that the special disks can hold up to 100 MB of data. As in the case of Iomega's Zip drives, an external drive connects to the computer's central processing unit (CPU), or an internal drive resides inside the CPU next to the floppy drive. For more information about Iomega products, view them on the Web at http://www.iomega.com. Unlike ADSM, removable media is a manual backup process, but it is user friendly and priced under $200 (prices vary per model).

Prevention is the Key

As a computer user, it is imperative that you back up data regularly as insurance against unpredictable computer problems. If you have a system administrator, contact this individual about the best backup media for you. If you do not have a system administrator, e-mail the IA Project's Desktop team (ia-deskcore@lanl.gov) for recommendations about backup tools that match your computer-skills comfort level. Lastly, if you do have a backup method in place, remember to back up your important data regularly and verify that the backed-up data can be easily and fully restored.

Electronic Information Protection Regimes

by Tad Lane, Information Architecture Standards Editor, CIC-1 Communication Arts and Services

"This guidance should clear up the widespread confusion over ... the differences between similarly named and seemingly boundless categories such as 'unclassified controlled nuclear information' and 'sensitive but unclassified nuclear information.'"

- "Science at its Best, Security at its Worst", PFIAB, June 1999

Question: What do the following all have in common—ECI, UCNI, NNPI, S&T, DUSA, AT, OUO, C/FGI-MOD?

Choices:

A. They're all unclassified.

B. Each overlaps with one or more of the others.

C. Taken in combination, they're only part of the realm of unclassified information that the Laboratory handles.

D. Most of us don't need to know what all the acronyms mean.

E. All of us do need to know how to handle whatever information we encounter in our job.

The answer, of course, is F, all of the above.

In the Information Architecture (IA) Project, we've been working on sorting out issues related to the handling of unclassified electronic information for a number of years now. With the help of Laboratory Counsel and various Laboratory security officers, we've gone back over applicable laws, orders, policies, guidelines, manuals, and drafts. Along the way, we've learned that there's an awful lot of information about information that, for most people, just isn't worth learning. It's too complicated, sometimes contradictory, and a mess—in spite of the fact that every element, when taken in isolation, has a clear and meaningful reason to be there.

For example, what is the difference between the CPRA and FOIA? There is an answer, but it's not the type of thing most of us really want to know. (Okay, the Laboratory is bound by the CPRA (California Public Records Act), while DOE is bound by FOIA (the U.S. Freedom of Information Act), which means that when Joe Public asks for information about the Laboratory, the request falls under the CPRA if it goes directly to the Laboratory, but it falls under FOIA if it goes through DOE to the Laboratory. That's part of it, at least. There's also stuff like pending litigation being covered by the CPRA but not FOIA, etc. Not exactly the kind of thing you'd expect St. Peter to want you to know.)

Like anything that involves law, however, there is always a "however." In this case, it's that however inaccessible the definitions of the various flavors of controlled unclassified information may seem, we're still all required to handle that information appropriately. Bad press and lawsuits really aren't the problem. The real risks are violations of individuals' privacy, providing shortcuts to sometimes hostile entities developing their aggressive capabilities, providing terrorists with information they need to do things we don't want them to do, etc. All in all, good reasons to control the information.

To resolve this conflict between the complexity of information categories and the simple need to appropriately protect information, we in the IA worked out a model based on the protection that information requires, rather than the category of the information itself. We're not concerned about what the information is, only how to protect it. From this angle, we developed the Laboratory standard IA-6303: Electronic Information Protection Regimes, which spells out four basic levels for protecting unclassified Laboratory information:

- Protection Regime 3: "Robust Authentication, Authorization, and Encryption"
- Protection Regime 2: "Robust Authentication and Authorization"
- Protection Regime 1: "Simple Authentication"
- Protection Regime 0: "Unrestricted Dissemination"

Acronyms and Such

ADC – Authorized Derivative Classifier

AT – Applied Technology

C/FGI-MOD – Confidential Foreign Government Information-Modified

CIPA – California Information Practices Act

CPRA – California Public Records Act

CRADA – Cooperative Research and Development Agreement

DES – Data Encryption Standard

DUSA – Designated Unclassified Subject Area

EAS – Employee Authorization System

ECI – Export Controlled Information

FOIA – U.S. Freedom of Information Act

FOUO – For Official Use Only

LA-UR – Los Alamos Unclassified Report

NIST – National Institute of Standards and Technology

NNPI – Naval Nuclear Propulsion Information

OSTI – DOE Office of Scientific and Technical Information

OUO – Official Use Only

S&T – Scientific and Technical Information

SSL — Secure Socket Layer

SUTI – Sensitive Unclassified Technical Information

TSPA – Technology/Software Publicly Available

UCNI – Unclassified Controlled Nuclear Information

Within this model, we start at the top. Does information require the highest level of protection? If not, does it require the next highest?

Some Basics on Unclassified Information

Before going much further into the model, it may be helpful to go over a few basic concepts involving the treatment of unclassified information.

1. There will always be a conflict between the public's right to know and our requirements to protect information. The public pays for the work we do; the public ultimately owns the infor-mation we generate; and the public is, in general, entitled to know what that information is. At the same time, though, the public, through its representatives in government, has established laws, orders, etc., which place limits on our ability to disseminate information. There are conflicting pressures here, we can't change that, and we just need to accept that as our environment.
2. It is the content owner's responsibility to determine what level of protection information needs. No one of us can or should be expected to understand all the different types of information at the Laboratory. But all of us need to understand the type of information we generate.
3. It is the content owner's responsibility to label information so that others can know the level of protection it needs. We can't know what we're not told. Labels can be attached to electronic information as coversheets or as banners at the top and bottom of pages.
4. It is everyone's responsibility to follow the labeling that the content owner has attached to information. If the content owner says that something requires robust authentication (i.e., SecurID or CryptoCard passcode), then we should protect it to that level.
5. It is everyone's responsibility to ask questions when they arise. If we have reason to think that something may require more protection than its label calls for, then ask the content owner. If we have unlabeled information and don't know what protection it requires, then ask the content owner.
6. All information is sensitive. To one degree or another, all information at the Laboratory requires some level of protection. If it's available to the public, for example, we need to protect it against hacking so that what the public reads is in fact accurate.
7. Information should be protected according to its highest level of sensitivity. For example, information that is "Official Use Only" (OUO) might also be "Unclassified Controlled Nuclear Information" (UCNI). OUO only requires minimal protections in and of itself, but UCNI requires very high protections. Information that is both OUO and UCNI should be protected to the level that UCNI requires.
8. None of this has anything to do with classified information. The protections for unclassified information sometimes parallel the protections for classified, but that does not mean that classified information belongs on unclassified computers. Ever. Whenever we have any questions about whether information might be classified, we need to check with S-7, Classification, or our Authorized Derivative Classifier (ADC).

9. Everything is subject to change. We make our best efforts to operate within the current technical and regulatory environment, but the environment and its requirements are continually evolving. As examples:

    a. New DOE policies currently under consideration may replace the OUO designation with FOUO ("For Official Use Only") and may bring some changes to what the designation means and how it should be used. Until or unless FOUO takes effect, however, we need to continue using OUO as outlined below.

    b. A new SUTI ("Sensitive Unclassified Technical Information") is currently under consideration which may require stronger protections for certain types of unclassified technical information. If this change does take effect, it will narrow the range of information that the protection regimes apply to, but it will not affect the protection regimes themselves.

Given the above, there are three basic components to information protection:

- confidentiality, controlling who gets to see information;
- integrity, ensuring that information is accurate; and
- availability, ensuring that information can be accessed when it is needed.

The requirements for integrity and availability vary on a case-by-case basis, depending on the value of the information and the audience that needs to access it. Integrity and availability requirements for a lunch date are generally pretty low. Integrity and availability requirements are high, however, for hazardous materials information that emergency workers need to access in case of fire. Similarly, integrity requirements are very high for information that is available to the public and that represents official Laboratory policy. (Another component of availability is accessibility, ensuring that information can be understood when it is accessed.)

Unlike the case-by-case approach required for integrity and availability, confidentiality requirements can be grouped into basic levels. It is confidentiality that the IA addresses with its protection regimes.

The Four Regimes

As a final bit of background, we chose the word "regime" to describe protection levels because it distinguishes the protection from the information itself and it doesn't carry the implications that some other words have. For example, if we had said "protection classifications," then that would have gotten tangled up with classified information, which is something we don't want to do. "Regime" is just a clean, unencumbered word which can describe levels of protection that multiple types of information can share.

Note: These protection regimes focus on accessing information from servers. There are still issues to be resolved regarding the protection of information on desktop client machines. For now, the general recommendation for desktop machines is to "use good practices" (start-up passwords, password-protected screensavers, etc.). More specific guidance may emerge in coming months.

Protection Regime 3: Robust Authentication, Authorization, and Encryption

This is the highest level of protection for unclassified information in the Laboratory standard. "Robust Authentication" means a SecurID or CryptoCard passcode, or the equivalent. "Authorization" means that the individual who is given access has specifically been given permission to access the information based on a need to know. "Encryption" means that any transmission of the information across the open network needs to be encrypted.

In more specific terms, authorization is generally based on permissions in the Employee Authorization System (EAS), and encryption for Web documents generally means Secure Socket Layer (SSL) using the Data Encryption Standard (DES) or Triple-DES. In some cases, such as UCNI, DOE requires that the encrypting software be certified by the National Institute of Standards and Technology (NIST). Encryption is not generally required for storage or for transmission across a shared Local Area Network (LAN), though there may be cases where it's desirable. Also, for all of this, "equivalent protection" is acceptable if it has been approved by S-5, Computer Security.

Information that requires this level of protection is generally information that we are required by law to strongly protect. Typical examples include the following:

- UCNI (Unclassified Controlled Nuclear Information), which is unclassified information that might significantly increase "the likelihood of (a) illegal production of nuclear weapons or (b) theft, diversions, or sabotage of nuclear materials, equipment, or facilities," and which we are required to protect under the U.S. Atomic Energy Act.
- C/FGI-MOD (Confidential Foreign Government Information-Modified), which is information from foreign governments that is not considered classified by DOE but which DOE has agreed to strongly protect.
- AT (Applied Technology), which is information designated by DOE that deals with technology advances in areas with major funding emphasis.

Protection Regime 2: Robust Authentication and Authorization

This regime has all the protections of Protection Regime 3, except that encrypted transfer is optional instead of required (though recommended where appropriate). Access is still restricted to individuals with an authorized need to know, as established through robust authentication (SecurID or CryptoCard passcode, or the equivalent) and EAS-based authorization (or the equivalent).

Again, information protected at this level is generally information that we are required by law to restrict access to. Examples include the following:

- ECI (Export Controlled Information), which is technical information that does not qualify for export under the unrestricted Technology/Software Publicly Available (TSPA) export label, which in turn generally means (a) unclassified information "related to nuclear weapons or nuclear energy technology, conventional weapons, or advanced technologies which could be of assistance in developing such controlled items," or (b) advanced cryptology tools and techniques.
- Personal/Medical information, including employees' medical, personnel, and security records, which are protected under the U.S. Privacy Act, the California Information Practices Act (CIPA), and University of California policy.
- Protected Cooperative Research and Development Agreement (CRADA) information, which is information we have agreed to protect as part of the agreement with our CRADA partner.
- Procurement activities, including information related to proposals, selection process details, and other procurement activities prior to official announcement of the selections.

Protection Regime 1: Simple Authentication

This regime requires only "simple authentication," which refers to fairly weak screening to establish that a user is authorized to access Laboratory internal information (e.g., an employee, a contractor, or a collaborator). Techniques for establishing this include IP address screening, simple Web passwords, or simply placing the information in the Unclassified Protected "Blue" Network.

Information with this level of protection is generally information that we are not required to but are given the discretion to restrict access to. Such information would not significantly harm the Laboratory's interests if it were publicly released, but it does not represent the Laboratory's "official public presence." Examples include the following:

- Drafts or works in progress that we intend to eventually release to the public but which have not yet been completed.
- Material marked OUO for discretionary reasons, which does not include OUO that falls under more restrictive protections. (More on OUO below.)
- Laboratory-internal discussions where we're expressing personal opinions that do not represent official Laboratory policy.

Protection Regime 0: Unrestricted Dissemination

This regime is for information that can be freely disseminated throughout the Laboratory and the external Internet. There are no restrictions on read access, though integrity and availability should still be maintained, commensurate with the value of the information, in order to ensure its accuracy. For Web postings, note that it is important to use the Laboratory standard copyright (to protect the public's legal right to access the information) and disclaimer (to protect the Laboratory's legal interests).

Information that falls under this regime includes (a) information that has been formally released through a process such as the Los Alamos Unclassified Report (LA-UR) release process, and (b) information that is not technical or scientific, is not sensitive, has no commercial value, is not a "publication," and otherwise does not require access restrictions. Examples include the following:

- The Laboratory Web home page, which does not include sensitive information and is not a "publication."
- Laboratory news releases, which are released through Public Affairs.
- The IA Project Web home page, which is covered by a DOE-approved Designated Unclassified Subject Area (DUSA).
- LA-UR reports, which have been approved by S-7, Classification, for public release.

Some Notes on OUO

OUO is a widely used designation for information that requires some level of access protection. It's a popular designation, but there are risks and problems associated with it.

For starters, OUO has a specific DOE definition, which makes it a lot less "fuzzy" than designations such as "in confidence." OUO is defined by DOE as information which falls into one of a series of "exceptions" to the FOIA and which meets a "sensitivity test." The "sensitivity test" is whether, in the opinion of the originator of the information, the Government's legitimate interests in restricting access to the information outweighs the public's right to know. But that sensitivity test can only be applied to information that falls into one of the specific exclusions to the act, including

- information that is protected by Executive Order or other statute (UCNI, Privacy Act, etc.);
- information that is related solely to internal personnel rules and practices;
- trade secrets, proprietary information, etc., from outside sources;
- communications that regard matters which are not yet official policy or which express personal opinions regarding policy;
- personnel and medical files; and
- law enforcement, banking regulation, and geophysical well information.

The California Public Records Act (CPRA) has a similar set of exceptions, with pending litigation added. (As previously mentioned, FOIA applies to the DOE, while CPRA applies to the Laboratory.
The similarities between the two acts show that, in practical terms, this distinction has little effect for most of us.)

As the list of exceptions shows, there is a wide range of information that can be labeled OUO, and there are a wide range of protections that OUO information can require. At the one extreme, there is information such as UCNI, which requires Protection Regime 3. At the other extreme, I might send you an e-mail where I take a swipe at some Laboratory policy, which wouldn't warrant anything more than Protection Regime 1.

The fact that it is so broad weakens the meaning of OUO as a label. If I give you a file that is marked simply OUO, you have no immediate way of knowing what level of protection it requires. Should you password protect it? Can you share it freely among your associates? (Or, to get to the silly extreme, is it marked OUO because it's an exception to FOIA, the CPRA, or both?)

Hence, whenever possible, my personal advice is to use a more specific designation. If information is both OUO and UCNI, mark it UCNI, since that's the more restrictive category. If it is both OUO and a protected CRADA trade secret, then mark it as a protected CRADA trade secret. This can help deter inadvertent but inappropriate dissemination of information.

Also, as detailed below under "Page Markings," it can help to mark the level of protection that is required, instead of just the fact that the information is sensitive. For example, "Protection Regime 3: Robust Authentication, Authorization, and Encryption" is more descriptive than simply "UCNI."

Note: As previously mentioned, DOE is considering replacing the OUO designation with FOUO ("For Official Use Only"). If the change takes effect, then FOUO might have somewhat different protection requirements than OUO. The main points of the previous paragraphs remain un-changed, however—protect information to its highest level of sensitivity, and make markings as specific as possible.

Blue or Green? Where to Put the Information

With the implementation of the new unclassified network firewall earlier this year, we now have an Unclassified Protected Network (informally, "blue"), intended primarily for internal use, and an Unclassified Open Network (informally, "green"), intended primarily for communications with external collaborators and the public at large. In general terms, then, information intended for Laboratory-internal use belongs in the Unclassified Protected Network, while information intended for external use belongs in the Unclassified Open Network.

Having said that, the firewall was not really built to protect information. Instead, it was primarily built to protect our internal network resources against unauthorized attacks. Hence, for all of the protection regimes, where to put the information depends on who will be needing to access it.

For Protection Regime 3: Robust Authentication, Authorization, and Encryption, and Protection Regime 2: Robust Authentication and Authorization, the information almost always belongs in the Unclassified Protected Network. Note, however, that the firewall alone does not provide adequate protection for these regimes, since there is a wide mix of people who have authorized access to the Unclassified Protected Network. Instead, these regimes require SecurID or CryptoCard passcode protection, access authorization, and, when applicable, encryption, in addition to the protections offered by the firewall. In certain cases where external users need access to the information but do not have access to the Unclassified Protected Network, information in these regimes may be placed in the Unclassified Open Network, but check with S-5 to work out an approved plan for providing the required protections.

For Protection Regime 1: Simple Authentication, either the Unclassified Protected or Unclassified Open Network is acceptable, depending on who needs to access it. For an audience that is limited to authorized users of the Laboratory network, the Unclassified Protected Network makes more sense, since the firewall already provides the simple level of authentication needed to meet the regime's requirements (though additional protections can be added if appropriate). For an audience that includes external users who do not have access to the internal network, then the information can be placed in the Unclassified Open Network with simple password protection or IP address screening. Note, however, that this only applies to read access for confidentiality protections; write access for integrity and availability still needs stronger protections.

For Protection Regime 0: Unrestricted Dissemination, the Unclassified Open Network is frequently the best place to put the information, especially if the information is intended for a public audience. If the information is intended for a purely Laboratory internal audience, then it is acceptable to place it in the Unclassified Protected Network. Before posting information for public access, remember to work with S-7, Classification, for whatever publication release might be needed (e.g., if it would be considered a "publication" in printed form, then it needs publication release in order to correctly record it and attribute it to the Laboratory for OSTI and Appendix F). And again, protect write access for integrity and availability.

Page Markings

As previously mentioned, it is the information originator's responsibility to determine which protection regime information requires. If it is scientific or technical information, work with your ADC or S-7 to ensure that it is not classified before placing it anywhere on the unclassified network. If it is something you are writing for others, then work with the content owner to determine the required level of protection.

After determining the protection regime, the information originator needs to communicate that to whoever receives the information, so that appropriate protections can be maintained. In informal settings among small teams, this might be communicated by word of mouth. For cases where there is a single point of entrance to access the information, such as through a required login screen, a banner page with appropriate notifications can do the job.

For basic Web pages, the Laboratory standard is IA-6501: Standard Page Markings for the Laboratory WWW. In addition to markings such as the standard disclaimer and copyright notice, this standard specifies the following markings to indicate the level of required protection:

- For Protection Regime 0: Unrestricted Dissemination, protection markings are optional. If desired, "Unlimited release" can be added to the bottom of the page to let the reader know that confidentiality need not be maintained for the information. If an LA-UR number has been assigned to a Web publication, then that number should be shown on the top-level page (e.g., title page) and at the bottom of all included pages.
- For Protection Regime 1: Simple Authentication, a marking at the bottom of the page can notify the reader of the intended restrictions. For example, "Laboratory internal use" or "Government use only" both indicate an audience that the information should be limited to.
- For Protection Regimes 2 and 3, prominent notices should be placed at the top and bottom of each page that indicate the protection regime, the type of information (UCNI, CRADA, etc.), and a description of the restrictions on its distribution. The following is the example from the standard:

Protection Regime 3: UCNI

This is Unclassified Controlled Nuclear Information that requires authentication, authorization, and encryption protections.

For More Information

A full list of the Laboratory standards discussed in this article can be found in the IA Web Standards Summary at http://www.lanl.gov/projects/ia/summary/web.html . Further discussion about related issues can be found in the following BITS articles: "Web Security in the Open Network Security Model" (April 1997), "Responsible Use of the Internet" (July 1996), and "Copyright and the World Wide Web" (May 1996). For more information about the IA Project in general, visit our project home page at http://www.lanl.gov/projects/ia/ .

Web Content Accessibility: New W3C Guidelines Have Arrived

by Tad Lane, Information Architecture Standards Editor, CIC-1 Communication Arts and Services

Hwæt we Gar-Dena in geardagum
þeodcyninga þrym gefrunon,
hu ða æþelingas ellen fremedon!

—Beowulf

Each of you reading this knows things that I do not. A language I've never learned, a movie I've never seen, a deeper insight into the Lorentz transformations. Whatever it is, there are things that each of you could say that would completely baffle me.

Meanwhile, I'm betting that I can return the favor by baffling most of you with the above quote. Yes, I know what it means. Yes, I put it in a place where it is available for you to read it. And yes, I deliberately put it into a form that is available but incomprehensible to most of the BITS audience. I made it inaccessible.

Just to get it out of the way, the Donaldson prose translation of the start of Beowulf is "Yes, we have heard of the glory of the Spear-Danes' kings in the old days—how the princes of that people did brave deeds." The point of it has nothing to do with what the words mean, however. The point is that I made the words available to you (either electronically or in print), but unless you happen to know Old English, I made things inaccessible and our act of communicating broke down.

Ever since the first primitive graphics began appearing on the Web, Web content accessibility has been an issue. Web designers' well intentioned desires to make pages attractive have too often led to pages that are unreadable for large segments of Web users. As advances in scripting and active content have increased Web designers' capabilities, the issue has only become more important.

Fortunately, as our capabilities for making content inaccessible have grown, so has our understanding of techniques for counteracting that inaccessibility. Back in 1997, the World Wide Web Consortium (W3C) launched the Web Accessibility Initiative (WAI). A number of resources have emerged from the WAI, leading up to this year's W3C Recommendation, "Web Content Accessibility Guidelines 1.0."

This Recommendation, in combination with supporting documents dealing with issues such as user agents and techniques, provides the most thorough set of good, common sense accessibility guidelines I have yet read. It prioritizes its guidelines so that we can distinguish between Priority 1, which we must meet or else it will be impossible for certain groups to access our information; Priority 2, which we should meet or else it will be difficult for certain groups to access our information; and Priority 3, which we may meet in order to improve accessibility. It provides specific, prioritized checkpoints for us to measure our Web pages against. And it makes sense.

As of this writing (July 1999), the Information Architecture Project has proposed a Laboratory standard that "endorses the Web Content Accessibility Guidelines" and "encourages Laboratory Web authors and publishers to strive for the highest level of conformance practicable." This proposal, IA-9701: Web Content Accessibility, may have been approved as a standard by the time this article is published. It can be accessed at http://www.lanl.gov/projects/ia/stds/ia970110.html.

In this article I'll present a high-level overview of some of the guidelines and themes in the Recommendation and its supporting documents. For more detail, I strongly encourage readers to go to the W3C Web site to read the Recommendation itself (URL provided under "For Additional Information", below).

We're Not Just Doing it for "Them"

An ongoing misperception is that accessibility only matters in "special" cases where we're worried about access for sight or hearing impaired users. The first bad implication of this idea is that a large group of users somehow do not have the same fundamental right to information that others have. Are they supposed to settle for only those scraps we choose to make accessible?

The second bad implication is that sight and hearing impairments are the only impairments we're worried about. That's not true. Among others, we also need to pay attention to users

- who may have difficulty moving a mouse,
- who may be using a different browser,
- who may have difficulties with certain color combinations,
- who may have chosen to turn off their graphics for performance reasons,
- who may have slow network connections.

And while we're at it, we ought to be paying attention to user agents that are not people at all, such as indexing robots.

Accessibility enhancements have a tendency to provide unexpected benefits beyond the ostensibly targeted group. The classic example here is the ramps that were put into the corners of street curbs, largely in response to the Americans with Disabilities Act. Although ostensibly targeted at wheelchair users, these ramps have also proven useful to bicycle riders, people pushing baby strollers, skateboarders, and others. Accessibility benefits us all.

I seemed to get a pretty good laugh at the IntraLab 99 Web Conference when I suggested the notion of browsing the Web while we're driving to work. After all, with the exception of those boneheads who read the Journal while driving between Albuquerque and Santa Fe, our eyes are occupied with the road, and our hands, with steering and shifting. But give us a voice browser and a wireless Internet connection, and browsing while driving can become as common as listening to National Public Radio. For people like me who have long, unproductive commutes, that possibility holds value. It will only work, though, for content that has been made accessible.

Separate Structure from Presentation

One of the broad themes highlighted in the W3C Techniques for Web Content Accessibility Guidelines 1.0 is that the structure of content should be distinguished from its presentation. In the pure model, each document has a Document Type Definition (DTD), which defines what each element is (a header, a list item, etc.), plus a style sheet, which defines how to present it (bold type, right aligned, etc.). In this model, someone who is using a voice browser can substitute their own style sheet so that the content can be more understandably presented in the voice environment.

Of course, the Web grew up too fast. Responding to market pressures, browser manufacturers opted for shortcuts (such as the <BIG> tag, which puts presentation into the middle of the document instead of in the separate style sheet). With the application of visual page design concepts to the electronic environment, Web designers used HTML markup for presentation instead of structure.

A few years ago, a common abuse of markup was to use heading tags to control font size, as opposed to correctly using heading tags to indicate headings. For example, <h3> might be used to designate bold, kind of big lettering, while <h2> would be for bigger bold lettering. Nowadays, I'm seeing more examples of the opposite abuse: tags being used to control presentation for headings, and the heading tag itself being omitted.

Consider, though, how an indexing robot or a voice browser which understands markup might read the document. The robot might be looking for the heading tags so that it can present an outline structure of the document to its users, a table of contents, if you will. The voice synthesizer might say "heading, level two," when it encounters an <h2> in order to communicate to its user the structural meaning. Unless the heading tags are correctly and consistently used to designate structural meaning, neither of the above tools will be able to correctly interpret the document.

Hence, the tags themselves should be used to signify what the element is (a level two heading, a list item, etc.). For presentation, as much as possible, use style sheets. (I'm planning a BITS article on style sheets for the near future. In the meantime, the best advice I can offer is to test under multiple browsers and to look for rendering that is acceptable and understandable, rather than ideal, under older browsers.) If style sheets do not yet provide the effect you're looking for, then try tag attributes (such as <h2 align="center"> for a centered heading). Finally, if all else fails but the presentation effect is still important to you, then use tags such as the tag within the structural tags so that the structure is still retained (such as <h2>Heading Text</h2>).

Provide Consistent, Logical Navigation

This is another example of how attention to accessibility can benefit all of us. By placing navigation tools (such as navigation bars) in a consistent location on each page within a site, all users can better understand both where to find the tools when they want them and how to bypass the tools when they'd rather get to the content more quickly.

In order to extend the accessibility of these navigation tools, do not rely on graphics alone to communicate them. Instead, if graphical navigation bars are used, then provide text equivalents that non-graphical user agents can follow.

- If the navigation bar is made up of a number of separate images, with each image linking to a different place, then provide an *alt ="description"* attribute for each image.
- If the navigation bar is an image map, then, in general, prefer client-side image maps to server-side image maps, since the client-side image maps allow you to specify an alt description for each region. (Also, do include the alt descriptions.)
- If the image map needs to be server side (due, for example, to unusual geometric shapes), then provide equivalent text links in some other accessible area.

Validate the Content

As the W3C techniques document points out, validating our Web pages is not, in and of itself, enough to ensure accessibility. Invalid or non-standard pages are, however, sure to reduce accessibility, not only for users of special tools but also for anyone using a browser that does not happen to support the particular bug and/or "feature" that might make the page look acceptable on your own browser.

Among the types of validation that the W3C recommends are the following:

- Spell and grammar checking. A misspelled word is bad enough for sighted users who can attempt to decipher it in context. For a voice synthesizer or an indexing robot, the misspelling can make it impossible to correctly process the word.
- Markup and style sheet validation. The W3C offers validation services that can test HTML markup and style sheets against the W3C standards. These can help catch errors which your own browser might compensate for but which other user agents may not (such as invalid nesting of tags).
- Automated accessibility validation. Tools are emerging which can test pages for known accessibility barriers. These can help catch common barriers, but they are not alone adequate to fully ensure accessibility (e.g., they can make sure that each image has an alt description, but they cannot ensure that the description makes any sense).
- Test with multiple user agents. Try out different browsers with different settings. Turn off graphics and see what the page looks like. Try to navigate the page without using a mouse.

Links to validation tools mentioned above are available from the W3C techniques document.

The Basic W3C Accessibility Guidelines

I'll wrap this article up with a list of the high-level guidelines included in the W3C Recommendation. There are further details, including specific checkpoints, in the Recommendation itself.

1. Provide equivalent alternatives to auditory and visual content.
2. Don't rely on color alone.
3. Use markup and style sheets and do so properly.
4. Clarify natural language usage.
5. Create tables that transform gracefully.
6. Ensure that pages featuring new technologies transform gracefully.
7. Ensure user control of time-sensitive content changes.
8. Ensure direct accessibility of embedded user interfaces.
9. Design for device-independence.
10. Use interim solutions.
11. Use W3C technologies and guidelines.
12. Provide context and orientation information.
13. Provide clear navigation mechanisms.
14. Ensure that documents are clear and simple.

A Final Admission

In the interests of fairness in advertising, I don't really know much Old English. I just had to memorize the first 11 lines of Beowulf back in college, and along with that I wound up remembering what those lines mean.

For Further Information

The W3C Web Content Accessibility Guidelines 1.0 is available online, along with links to its supporting documents, at http://www.w3.org/TR/WAI-WEBCONTENT/ . I highly recommend reading it and testing our own pages against its checkpoints.

The following BITS articles have dealt with related issues:

- "Making the Web Accessible Part 2: Text Issues and Conclusion" (December 1997)
- "Making the Web Accessible Part 1: Overview and Graphics" (November 1997)
- "Tips on Writing HTML <TABLE>s" (February 1996)
- "Why Not <BLINK> and <CENTER>? Writing HTML for Portability" (November 1995)

For more information about the work the IA Project has been doing with the Web, please visit our Web team home page at http://www.lanl.gov/projects/ia-lanl/area/web/ . For more information about the IA Project in general, visit our project home page at http://www.lanl.gov/projects/ia/ .

Lecture Review: Next-Generation Chips, Processors, Transistors, and Wiring

by Kimberlyn Mousseau, Technical Staff Member CIC-15, Advanced Information and Database Technology

The future of the computer chip was the focus of a recent lecture by Dr. Randall D. Isaac, Vice President of System Technology and Science at IBM's T. J. Watson Research Center. Dr. Isaac spoke March 30, 1999 at the Director's Colloquium at Los Alamos National Laboratory (LANL). He received his Ph.D. from the University of Illinois and has worked for IBM for the past 20 years in its Technology Development Department in Yorktown Heights. Working both on the design and production side of process and memory chips, he now has worldwide responsibility for strategic research in transistors and display devices for IBM.

In his presentation, Dr. Isaac discussed both present-day chip technology and future expectations. Addressing LANL, Dr. Isaac described his pleasure speaking about the future of computing at LANL, a site renowned for its history of computing. Convinced LANL will play a major role in the future of computing, he looks forward to close cooperation with many LANL scientists as they chart a future for computing. Dr. Isaac discussed current product-development plans, the history of the computer chip, and future developments.

Moore's Law and Gigahertz and Gigabit Chips

Product-development plans are currently in place for both gigahertz and gigabit chips, both of which are expected to be released in 2001. According to Dr. Isaac, it is important now to focus on where we go beyond these points and the technology challenges for moving beyond them. However, before looking into the future, he stated that it is essential to understand the history associated with gigabit and gigahertz as well as Moore's law.

What is known as Moore's Law, proposed by Gordon Moore, was first articulated in 1965 and published in *Electronics Magazine.* Moore predicted that the number of components per computer chip would double every year.

A decade later, in 1975, he published another paper in the *"Proceedings of IEDM."* It stated that the trend had continued, but would actually double each time in subsequent periods of no more than two years in length. Although Moore further predicted that the change in slope would occur in 1980, he was wrong—at least partially. As he mentioned in a 1995 article, the slope did change, but not in 1980, rather 5 years earlier than had been forecast.

Ultimately what was gathered from this information was that the slope isn't really what's important. Instead, what is important is Moore's statements about the cause and effect of this trend because, the historic trend enables us to better understand how to predict the future. The actual driving force behind the number of components per chip can be defined by a:

- 50% gain in increase of components per chip due to lithography resolution,
- 25% gain from device and circuit innovation, and
- 25% gain from increased chip size.

It is noteworthy that although Moore outlined these three factors in 1965, they remain true today. In addition, one of Moore's very important and profound points concerned the effect of the trend. He theorized that as function and speed increase, the cost per chip decreases. Another important factor, an extrapolation of Moore's Law that must be taken into consideration is performance improvement, which includes:

- Transistor performance,
- Interconnect density and delay,
- Packaging and cooling, and
- Circuit- and system-level gains.

Now let's talk about where we are today. It was first estimated that by the year 2010, a 64-gigabit chip would exist. Instead however, we will actually realize an 8-gigabit chip. In fact, a little bit of this trend is already beginning to surface—the 256-megabit is coming out about 9 months later than was previously expected. Dr. Isaac stated that he believes this trend will continue to slow where we will continue to see a change in the chip performance slope.

Chips Grow, but Processor Size Shrinks

As compared to chips, the future of what's in store for processor speed is more complex. If there were no further innovation, which is unlikely, we would only have an 8-gigahertz chip by 2010. And, if we stay on the predicted curve for technological advancements, we would have a 20-gigahertz processor by that same year. A 20-gigahertz processor implies that within a single cycle, you could not use the result produced from one calculation as input to another calculation at the opposite end of the chip. As we begin to approach these limits, we are likely to face some interesting challenges.

Another interesting contrast is that chip size may be growing, but processor size is shrinking. With lithography density moving as fast as it is, the number of transistors on a chip is growing faster than the number of transistors in a processor. Since the processor as a unit itself is getting extraordinarily small, the real gain will be in the system's integration. In fact, there are six dimensions to this integration that lead to the improvement of microelectronics. These are:

1. Lithography,
2. Transistor,
3. Wiring,
4. Memory,
5. Circuits, and
6. Design.

Improving Lithography

According to Dr. Isaac, lithography is the most important driver for improving microelectronics. Consequently, if we're going to progress into the future we must understand lithography. For 1999, technology is ahead of schedule by at least one year. The simple reason for this advance is that deep UV lithography was far more robust and easier to manufacture than expected. It took a long time to get into place, but afterwards it was very employable. Also, as the drop in DRAM prices (starting in 1995) contributed to the progress, manufacturers suddenly were forced to put more emphasis on improving lithography, and less pressure was placed on other parameters such as chip size. In fact, the release of the 256-megahertz chip marks the first time that a DRAM generation is coming out with a smaller chip than the previous generation. However, this state of affairs will not continue. We will slow down, and then get back on track. Once this occurs, it becomes more difficult to make predictions.

Where will lithography go next? Optical lithography is a moving limit and has moved steadily beyond our predicted capabilities. Clearly, if we're going to make the 2010 projection, we'll need to move past where we are today. The range of features we wish to improve can be characterized by wavelengths and categorized into an optical group. What limits us is material properties. We need to design and develop new and improved optical lenses—the lack of refractive lensing capability beyond the current status will be difficult.

Lithography is the significant portion of the driving force behind all trend lines. At present, we don't have a clear picture as to how to improve lithography. The bulk of future investments will go toward extreme UV, but it will be difficult because it will take many resources and will be expensive. Economics probably will play more of a problematic role than physics.

Improving Transistor Performance

Before discussing how to improve transistor structures, we must first look at the history of transistors. The laws of transistor scaling include the source, drain, and gate. When lithography gets smaller, transistors must get smaller, too. Changes must be made to voltage, along with gate oxide, wire width, gate width, diffusion, and substrate. This recipe is one that has been followed throughout the history of scaling.

The most important attribute of transistors, and the one demanding the most concentration, is the thickness of gate oxide. Products are currently at 3 nanometers of gate thickness and are pushing to 2 nanometers. We're encountering a fundamental change because we cannot change gate oxide thickness forever. If not changed, the limit may be around 1 or 2 nanometers. This won't mean the end of progress as long as we begin to think about ways to overcome the limitations of scaling, such as developing a new transistor structure.

To develop a new transistor structure, emphasis currently is being placed on silicon-on-insulator (SOI) because of its projected performance gain and lower power. SOI technology, according to Dr. Isaac, is a step along the evolutionary path to a radical new transistor structure, which may actually be a means to a bigger end, even though changing the structure will not be easy.

Other ways of improving transistor performance is to devise alternatives to silicon, produce higher mobility materials, and improve the limitations of low temperature. Another consideration is to improve solid-state coolers. Therefore, we need to stimulate research in low-cost, modular, environment-friendly, and reliable cooling technology. Transistor performance can continue, but it will require much innovation on structure, cooling, and materials.

Other Considerations

Other considerations for improving the computer chip include wiring, memory, circuits, and design. Copper interconnect technology is crucial. Today, copper wiring is being shipped in products where hierarchical wiring is key. It is very important to design philosophy to have many short-range small wires and longer-range connections with larger fatter wires. These things together with the lower-dielectric constant insulators that are coming into play, may allow interconnections to be extensible to about 100 nanometers and below, although there are still some real challenges to overcome. Also once the dimensions of wires have been shrunk, we'll need a line barrier thickness less than 8 nanometers, metal effective resistivity less than 1.8, and inter-level insulators that are less than 1.5.

Improving Performance

Industry will reach the gigahertz and gigabit era in the year 2001. Optical lithography will be extended, while the transition to "next generation lithography" will be difficult. New transistor structures (such as SOI) will be needed to continue the rate of performance enhancement. Copper wiring, low-k dielectrics, and hierarchical wiring designs will enable effective interconnects. By the year 2010, we can expect 8-gigabit DRAM and 8-gigahertz processors. There is a great potential to create new functions on a chip.

Tecolote: An Object-Oriented Framework for Hydrodynamics Physics

by Kathleen S. Holian, Technical Staff Member CIC-12, Scientific Software Engineering Group

In today's world, the advances in physics models and numerical models that are used in hydrodynamics are proceeding at an increasingly rapid pace. In computing it seems that new computer architectures are being put out practically daily. Against this back-drop of mind-boggling change, developers of new physics and new numerical models are increasingly being asked to rapidly determine just how feasible their new models are. Clearly, they can work more efficiently if they are able to build on what others have done and are only required to implement those parts of a model which are new. This environment calls out for the development of physics frameworks.

Tecolote, an object-oriented framework written in C++, is designed for the development and implementation of a wide variety of hydrodynamics applications. It is also meant for the rapid development and testing of all kinds of models, numerical or physical, that are related to hydrodynamics.

We have initially implemented two types of hydrodynamics using Tecolote—an Eulerian code and a Lagrangian code. These are the two main types of hydrodynamics algorithms that are used here at the Laboratory although there are many different incarnations and hybrids of each type in use. In an Eulerian code, physical space is divided up into computational cells, and material flows through the computational mesh. This means that cells can contain more than one material. A Lagrangian calculation again divides up the physical space into computational cells, but the cells are attached to a given object. They grow, shrink, or distort as the object changes. The cells do not contain more than one material.

We have also implemented a number of physics models in the Tecolote framework, which work in conjunction with any of the numerical hydrodynamics algorithms. These are such things as equation of state (relationship between density, temperature, energy, and pressure), material strength (description of material deformation at temperatures below melt), and high-explosive burn.

We have just released the Eulerian code (called Conejo) to users who are now using it for real calculations. The Lagrangian code (which is based on work done by Ed Caramana in X-Division[1]) is still in the development stage. However, it was implemented in record time, as a result of the fact that many of the modules originally written for the Eulerian code could be reused. This means that only the core numerical Lagrangian algorithm had to be written to have a fully functional new type of hydrodynamics code. Also, it is worth pointing out that the framework is flexible enough to handle completely different types of algorithms.

Laying Out the Ground Rules

Before designing the framework we laid out the ground rules for the features we needed to have. We have kept these rules in mind throughout the design and implementation of Tecolote, and we feel that we have been successful in adhering to them.

First, we wanted to have an object-oriented framework. This means that the modules in the framework are as independent from each other as possible. In object-oriented code design, the addition of new modules should be easier and more trouble-free than procedural code.

We also wanted the computer code to look as much like the original physics equations as possible. We did not want the high-level physics coding in existing modules to change if, for example, new hydrodynamics modules requiring differing underlying mesh geometries, were added to the framework.

We wanted the applications to be able to work in one, two, or three dimensions. In addition, we needed to be able to represent mesh-wide variables on a variety of centerings. An example of this would be a staggered-mesh application that uses cell-centered state variables but vertex-centered velocities (as does our Eulerian application). Certain physics operators, such as the divergence, need to know the centering of a variable.

Another important and crucial requirement was portability. Considering the new computer architectures that have multiple processors (or multiple boxes, each containing multiple processors), we wanted the high-level physics coding not to be cluttered with provisions for message passing and load balancing. We felt that these should be done at a lower level and in such a way that they could be tuned for different architectures.

The POOMA Library

The Tecolote framework is built on the POOMA Library, which is also written in C++. POOMA provides Tecolote with fields similar to Fortran-90 arrays but with extra features. The POOMA fields automatically take care of message passing on platforms with multiple physical processors. But POOMA also has a scheme for more efficient memory storage and, in the future, will have provisions for automatic load balancing. POOMA fields can be laid on top of different types of meshes, not just the standard Cartesian meshes of Eulerian codes. This feature will allow

for future development of hydrodynamics code applications that are on more complex meshes. The POOMA library also takes care of correctly performing mathematical operations, such as divergence and gradient, for a given mesh geometry, although these can be overwritten if desired in the Tecolote part of the framework.

No Wasted Storage Space

The compressed storage of POOMA fields is based on the idea of "virtual nodes." Virtual nodes are subunits of the mesh over which an application is doing its simulation. When a field variable is constant over an entire virtual node, the storage for that field variable is collapsed down to just one value for the whole subunit, rather than one value for each point of the grid in that subunit. This is somewhat analogous to the compression that is regularly practiced on graphics files that contain a great number of repeated values.

With Eulerian hydrodynamics codes, it is a common occurrence to initialize a great deal of mesh as void, so that the materials in a problem may expand into the void. Often many zeroes are stored for a field variable, wasting lots of storage space. The POOMA fields solve this problem.

Another situation is that in which one has many materials in a problem. Each material may have a number of state variable fields such as density, energy, and pressure. These fields need to be represented on the entire mesh of the problem. However, each material may be found only in a small part of the mesh. Using the POOMA fields, these material-dependent fields can be collapsed down to storing only one value across an entire virtual node if the field is constant across that node. Additionally, POOMA automatically keeps track of the fields. For instance, as a material moves across a mesh, virtual nodes expand and compress automatically, depending on the state of the field.

Tecolote Physics Library

The most important thing about Tecolote is that it provides a framework in which to develop new hydrodynamics models and methods. To encourage code reuse not only in the computer science arena, but also in physics, we plan to have a wide variety of physical and numerical solution modules available to the developer. The developer of a new method or model will then only be required to spend his or her valuable human time on coding the new model of interest. Other models that might be required to actually run a problem, such as the equation of state, would be available from the existing library of physics.

So far, we have implemented all the numerical/physics modules needed for an Eulerian hydrodynamics code. And, in fact, this code has been released to users now. In addition, a pure Lagrangian numerical solution algorithm has also been coded, and is being tested.

In addition, we have a large body of physics modules for various types of material behavior: equations of state, material strength, and high-explosive burn. And we are adding more variations to each of these types of models as time goes on. Some are simple, but some are very sophisticated and newly developed.

Tecolote Computer Science Features

The Tecolote framework also has a number of computer science features that were designed with the idea of easing the work required for developers of new models. In addition, a next-generation version of the computer science part of the framework is being developed, although not yet integrated with the physics part. In a future BITS article we will explain these features. An example feature is the ability to use default operators for different types of mesh and dimensions, or to override the default with a different customized calculation.

Current Computational Capabilities

We have run some very large problems on the ASCI Blue Mountain machine, using both the Eulerian and Lagrangian codes. For example, 1024 processors on 8 boxes were used for a real, 66-million cell Eulerian calculation. Although we have not yet done so, we feel that on the current architecture it may be possible to run up to a half-billion-cell Eulerian calculation.

We will soon have a capability to easily develop, test, and use new models that are useful for hydrodynamics applications in an object-oriented framework. This framework is designed to be available on a wide variety of computing platforms, including the new platforms that have complex multiprocessor structures. We are also emphasizing the ability to be able to simulate very large and/or finely resolved problems.

The author wishes to acknowledge the project team: L.A. Ankeny, S.P. Clancy, W.H. Dorin, J.H. Hall, S.R. Lee, J.C. Marshall, G.R. McNamara, J.W. Painter, S.J. Sydoriak, and M.E. Zander

Reference

[1] *E.J. Caramana, D.E. Burton, M.J. Shashkov, P.P. Whalen, "The Construction of Compatible Hydrodynamics Algorithms Utilizing Conservation of Total Energy,"* Journal of Computational Physics *146, 227–262 (1998).*

New Electronic Resources Available at the Research Library

Los Alamos Patents Available Online

by Kathy Varjabedian, Databases Team Member, CIC-14, Research Library

Intellectual property is one of an organization's most important resources, and at Los Alamos National Laboratory this is keenly evident in patents. LANL staff have been getting U. S. patents since 1947, and the Research Library has obtained copies of all of them and converted the paper copies to PDF files.

These Los Alamos patents are now available through the Research Library's online catalog. To find patents in the online catalog, search for Title word "patent" and another word of your choice, such as title, author, or abstract word. Catalog records such as figure 1 below provide identifying information, abstract, and URL link to the PDF file.

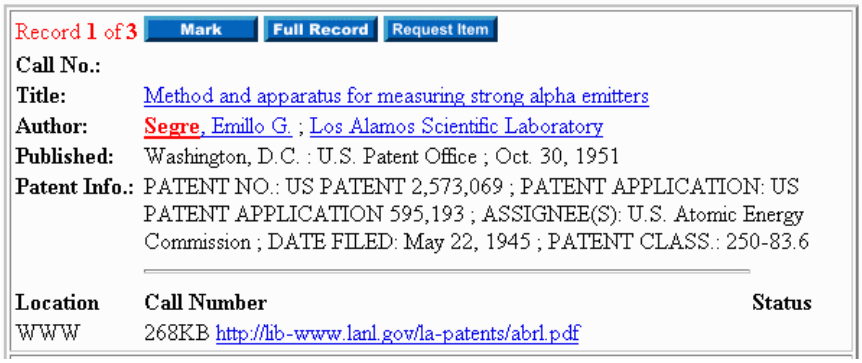
Record 1 of 3 Mark Full Record Request Item

Call No.:

Title: Method and apparatus for measuring strong alpha emitters

Author: Segre, Emilio G. ; Los Alamos Scientific Laboratory

Published: Washington, D.C. : U.S. Patent Office ; Oct. 30, 1951

Patent Info.: PATENT NO.: US PATENT 2,573,069 ; PATENT APPLICATION: US PATENT APPLICATION 595,193 ; ASSIGNEE(S): U.S. Atomic Energy Commission ; DATE FILED: May 22, 1945 ; PATENT CLASS.: 250-83.6

Location	Call Number	Status
WWW	268KB http://lib-www.lanl.gov/la-patents/abrl.pdf	

Fig. 1. Sample record from online catalog.

Analyzing the 1175 Los Alamos patents in the catalog shows that patents have increased significantly in recent years. Figure 2 shows the breakdown by decade is as follows:

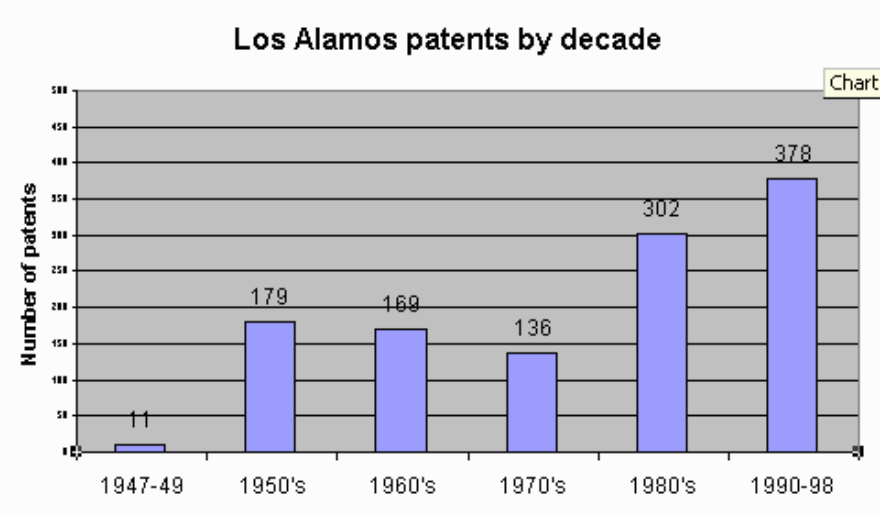

Fig. 2. Los Alamos patents by decade.

Figure 3 is an example of an interesting patent diagram:

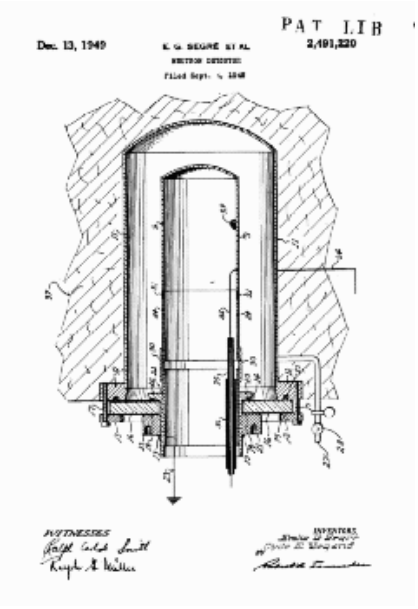

Fig. 3. U.S. Patent 2,481,222 Neutron detector, Emilio Segrè and Clyde Wiegand, 1949.

Free Patent Information on the Web

by Lou Pray, Customer Service Team Member, CIC-14, Research Library

European Patent Office offers free patent information on the web. The following URL is your online gateway to non-U.S. patents. http://ep.dips.org/ .

Searches are based on the bibliographic data in patent documents, using an easy-to-operate search interface. Use the form to specify the terms that you want to use in your search, and the fields in which these terms should be searched (such as inventor, title, IPC classification, etc.). Some full text images of the documents are available as PDF files.

The European Patent Office has broken up the database into the following files:

- European (EP) patents (most recent 24 months, with PDF files)
- PCT (WO) patents (most recent 24 months, with PDF files)
- Worldwide patents (coverage depends on country, mostly back to 1960s or 1970s)
- Japanese patents (1976-)

LOS ALAMOS NATIONAL LABORATORY

Research Library

http://lib-www.lanl.gov

The LANL Research Library offers a variety of training opportunities for the Laboratory community. Available sessions focus on specialized library databases and other electronic resources. A complete list of course offerings can be found at http://lib-www.lanl.gov/libinfo/training.htm. All sessions are available to individuals or groups, at the library or your site. Arrange for a session by contacting the Library, phone 7-4175 or e-mail library@lanl.gov. Library tours are available on a drop-in basis every Wednesday at 1:00 p.m.

Technical and Advanced Technical Computer Training Courses

Communications

- Eudora 4.x
- Meeting Maker 5.0.3

Office Skills 2000

- Office Skills 2000—LANL
- Office Skills 2000—Professional Development

Web Authoring and Browsing

- Dreamweaver 2.0–MAC or PC
- FrontPage 98
- HTML Basics
- HTML Intermediate

Enterprise Information Applications (EIA)

- Date Warehouse–Basics
- Date Warehouse–EDS Reports
- EDS—Basics
- EDS–GUI
- EDS–Training Plans
- Infomaker
- Invoice Approval System
- Purchase Card System
- Procurement Desktop
- Recharge
- Time & Effort GUI
- Travel Foreign GUI
- Travel Domestic GUI
- Web JIT

Other EIA Courses

- Financial Management Information System (FMIS)
- Property Accounting, Inventory and Reporting System (PAIRS)
- Signature Authority System (SAS)
- Secretarial/Contract Service (SE)
- Salary Review System (SRS)
- Directory Information System (DIS)
- Automated Chemical Information System (ACIS)

System Administration Training

- SGI System Administration (Beginning)
- SGI System Administration (Advanced)
- SGI Network Administration
- SGI Performance Evaluation and System Tuning
- Solaris 7 System Administration
- Solaris 7 Network Administration
- Solaris 7 Server Administration
- Unix and Widows NT Integration
- Windows NT Workstation and Server
- Windows NT Optimization and Troubleshooting
- Windows NT Security

Programming Training

- C Programming (Beginning)
- C Programming (Advanced)
- C++ for Experience C Programmers
- ANSI/ISO C++ Programming Clinic (Advanced C++)
- Distributed Objects Using Corba
- Java Programming
- Java Program Workshop
- Distributed Programming with Java
- Object Technology– A Management Overview
- Object-Oriented Analysis and Design
- Perl Programming
- Advanced Perl Programming with CGI
- C-Shell Programming
- Programming for Beginning using Java

Application Training

- Advanced WWW Development
- FrameMaker Basic and Advanced
- Foundations of IDL Programming
- IDL 5.0 Graphic Object Workshop
- Netscape Servers for Intranet Development
- Origin 2000 Applications Programming and Optimization
- Sendmail–Managing Internet Mail
- C++ and the Unified Modeling Language
- Sybase Fast Track to Adaptive Server Enterprise 11.5 (ASE)
- Sybase Performance and Tuning for System 11
- Sybase SQL Server Administration
- Unix (Beginning)
- Unix (Advanced)
- Visual Basic 5.0 Fundamentals
- Visual C++ Windows Programming

ASCI

- Mastering Projects Workshop
- Software Engineering for Scientists and Engineers
- Getting started on ASCI Blue Mountain Systems
- Running MPI on Blue Mountain Systems
- Introduction to Totalview
- LSF (Load Sharing Facility)
- Introduction to HPSS (High Performance Storage System)

** You do not need an ICN password to use e-mail.

Telling Your Story in BITS: Sharing Your Expertise and Energy

by Betty Katz, Technical Writer/Editor CIC-1, Communication Arts and Services

If you ever have a spare moment, please survey the folders on your unclassified computer desktop. How much of that information would help others learn a new technique, solve a technical problem, or find helpful resources to streamline their work?

BITS bets that you have a treasure of ideas that you would share happily with others verbally over a cup of coffee. If so, please give BITS a chance to explain how sharing your expertise in writing can actually be fun. What about highlighting some of this information in a monthly forum that spotlights the collaborative work of CIC Division with other Laboratory divisions? Consider the BITS audience—380 U.S. and foreign subscribers and almost 1,200 Los Alamos employees. What about writing an article for BITS?

Personal Energy is the Underlying Thread

It's easy to be energetic about sharing new information with a colleague or a family member. But what causes the inevitable power drain when someone suggests we write that story down for others to share? Ghosts of English teachers past appear, and creative excusery sets in. Suddenly we shift without using the clutch into a person totally different from the relaxed individual who told the story so naturally. No need for any BITS contributor to go into that mode. In fact, BITS is asking contributors to plug into their personal prowess as successful professionals and tell their stories.

BITS relies on you to share your ideas in written forms such as reports and articles. You don't have to submit polished writing samples to be considered for publication. Your job as a contributor is to muster the energy to get the information in writing. For each issue, the managing editor incorporates an editing and review cycle in the production schedule. But don't worry; you won't lose control of your story. You will get a chance to review any editing changes before your article goes to print.

If you're looking for incentives to get you going, consider BITS as a venue for getting a jump start on a conference paper or journal article you'd like to write. If improving your writing skills is a goal, consider your writing submissions to BITS as evidence for your performance appraisal.

Some Tips for Getting Started

Whatever motivates us to start telling our story, writing guru William Zinsser encourages us to be ourselves as we write. In his informal guide to nonfiction writing, *On Writing Well*, he offers authors the following advice:

> "Humanity is a fundamental principle of effective writing; be natural; be yourself when you write. Never write anything you would not say. In your writing sound like who you are. How you write is how you define yourself."

Following Zinsser's advice can help us avoid the well-worn pitfalls of robot expression that never uses anything when it can be utilized; that never ranks anything when it can be prioritized; that never motivates people when they can be incentivized.

If you would like additional writing guidance along the way, try tapping into the resources listed below. These authors understand what it takes to fill a page effectively. They also will remind you of the service you perform when you share new information clearly.

BITS staff members are convinced that you can fill the magazine's pages with news of your accomplishments. As a contributor, you will have the joy of reconnecting to those energy sources that stimulated your achievements in the first place. Your readers will appreciate this new information and will thank you for your efforts because as Zinsser reminds us, "Easy writing makes hard reading."

Resources for Writers

William Zinsser, *On Writing Well: The Classic Guide to Writing Nonfiction*, 6th ed. (Harper Perennial, New York, 1998.) (Also on audiocassette: Harper, Audio; ISBN 1559943491.)

Struggling to write becomes an energizing workout once the reader plugs into Zinsser's methods for trimming the fat from sentences and paragraphs. The author notes that "the secret of good writing is to strip every sentence to its cleanest components." To do this, Zinsser suggests that writers clear their heads of clutter. Clear thinking becomes clear writing.

But Zinsser admits, "A clear sentence is no accident. Very few sentences come out right the first time, or even the third time." He advises writers to look at their work and ask, "Is this clear to someone encountering the subject for the first time?" If not, the author suggests that some fuzz has "worked its way through the machinery." The clear writer," Zinsser affirms, " is someone clearheaded enough to see this stuff for what it is: fuzz."

William Zinsser, *Writing to Learn* (HarperCollins, New York, 1989).

Don't let the 1989 publishing date put you off. Check out the amazon.com Web site and read the five-star reviews for *Writing to Learn*. This book is an absolute joy for technical and scientific writers to read. With humor and understanding, Zinsser offers scientific writers commonsense support in their writing adventures.

Quotable insights fill the entire book; however, Chapter 6, "Earth, Sea and Sky," Chapter 8, "The Natural World," and Chapter 11, "Writing Physics and Chemistry," prove Zinsser's assertion that there is romance in technical and scientific writing. What about these supportive words for a struggling writer on a foggy afternoon?

> "Every science has its unique romance and a writer who can capture that romance enables the rest of us to glimpse what it is about the field that makes it exciting to those who love it. It's one of the best gifts that science writing can offer."

Hank Nuwer, *How to Write Like an Expert About Anything* (Writer's Digest Books, Cincinnati, Ohio, 1995).

This book's subtitle*, Bring Factual Accuracy and the Voice of Authority to Your Writing,* encourages writers to share their expertise with audiences hungry for reliable and up-to-the-minute information. Nuwer's advice underscores the need for writers to understand how an obsession with accuracy lends authority and authenticity to their writings. One of Nuwer's most practical tips is his suggestion that writers slant their piece so that everything from the first sentence to the closing paragraph can be reduced to a single, well-shaped idea.

Peter Elbow, *Writing with Power*, 2nd ed. (Oxford University Press, Inc., New York, 1998).

Professor Elbow suggests that reluctant writers adopt a sense of trust so that "when you have the germ of an idea or even just the hankering for one, you will be led sooner or later to the words you are looking for, if you just start writing." He doesn't mean that good writing just happens. He does mean that writing with power involves "getting power over yourself and the writing process; knowing what you are doing as you write; figuring out what you really mean…not feeling stuck or helpless or intimidated."

Ways to Contribute to BITS

If you'd like to contribute articles or ideas for articles, such as upcoming events or a project you'd like covered in BITS, contact the managing editor. To increase the content and quality of the publication, we'd prefer to cultivate an environment that is constantly bringing in words that can be used. People can contribute in several ways—essays, articles, reviews, interviews, team reports, and bits and pieces. You could also join the BITS Contributors' Board and help develop the publication. Contact the managing editor or send e-mail to the Contributors' Board listserv: bitsboard@lanl.gov.

Reader Feedback

Feedback helps us to provide a document that responds to the changing needs of its readership. If you have comments or questions about this publication, please let us hear from you. We have reserved the back of this form for that purpose. We also accept articles for publication that are of interest to our readers. Contact the managing editor for more information. This form is also used for new subscriptions, deletions, or changes. Instructions are on the back. If you prefer to contact us by e-mail, send your comments and/or subscription request to denise@lanl.gov.

Do Not Staple
Fold on This Line First

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL
FIRST-CLASS MAIL PERMIT NO. 88 LOS ALAMOS NM

POSTAGE WILL BE PAID BY THE ADDRESSEE

MAIL STOP B251
ATTN: DENISE SESSIONS, MANAGING EDITOR
CUSTOMER SERVICE GROUP (CIC-6)
LOS ALAMOS NATIONAL LABORATORY
PO BOX 1663
LOS ALAMOS, NM 87544-9916

Do Not Staple, Seal with Tape
Fold Here

cut along dashed line

Feedback

New Subscription, Deletions, and Changes

Bits is published by Los Alamos National Laboratory. If you would like to be added to or deleted from our mailing list, please check the appropriate line, complete the form below, and mail us the form, or e-mail bitsupdate@lanl.gov

_____ Add my name to the BITS mail list.

_____ Delete my name from the BITS mailing list.

_____ Change my name/address as indicated below.

Name _____ Date _____

Address _____ Mail Stop _____

Group _____ Organization _____

City _____ State _____ Zip _____

Phone _____ Number of copies _____

1999 12-Month Index

For a more complete index, see <http://www.lanl.gov/orgs/cic/cic6/bits/index/index_home.html

Keywords	Title of BITS Article	Date	Page
ASCI	*Blue Mountain IS World's Fastest Computer*	*Dec. 98*	*1*
ASCI Training	*ASCI Training Course Descriptions*	*J/J 99*	*27*
BITS Announcements	*BITS Gets New Design, New Editor*	*F/M 99*	*3*
BITS Interviews	*John Blaylock*	*A/S 98*	*2*
	Chris Kemper	*A/S 98*	*8*
	Bob Newell	*O/N 98*	*1*
	Earleen Eden	*O/N 98*	*3*
	Dale Land	*Dec. 98*	*11*
	John Morrison	*F/M 99*	*6*
	Ann Hayes	*J/J 99*	*22*
CIC (Computing, Information & Communications)	*BITS Welcomes New Deputy Division Director for Strategic Computing*	*F/M 99*	*6*
CIC-2	*CIC-2's Future Is in Stewardship, Automation, and "Fun Stuff"*	*Dec. 98*	*11*
CIC-10 CIC-14	*CIC-10 "Knowledge Is Our Most Important Product"*	*O/N 98*	*3*
Desktop Computing	*CIC-2's Future Is in Stewardship, Automation, and "Fun Stuff"*	*Dec. 98*	*11*
	New Ways for Managing Software Development Projects	*J/J 99*	*2*
	Scientific Software Process Engineering	*J/J 99*	*4*
Embedded Systems	*The Year 2000 Bug is Hiding*	*F/M 99*	*1*
ESD	*Working for and Serving You*	*F/M 99*	*10*
HTML	*Extensible HTML*	*J/J 99*	*12*
High Performance Computing	*Performance and Scalability Analysis of Apps. on T-flop-scale Architectures*	*J/J 99*	*7*
	Major 3-D Parallel Simulations	*J/J 99*	*9*

Keywords	Title of BITS Article	Date	Page
Infrastructure	*Virtual Private Network*	*J/J 99*	*15*
Information Systems	*Engineering Index Is Now Available from the Research Library*	*Dec. 98*	*13*
	Research Library Adds Social SciSearch® Database	*F/M 99*	*5*
	Library Databases	*J/J 99*	*20*
Newell, Bob	*24 Years Later, Bob Newell Joins the Lab as Deputy CIC Division Director*	*O/N 98*	*1*
Networks	*Laboratory Unclassified Network Will Implement Changes for Secuirty*	*Dec. 98*	*3*
	Preserving Bandwidth on Your Windows NT Network	*F/M 99*	*8*
PNG	*Its' Time for PNG: A Graphics Format You Can Pronounce*	*O/N 98*	*7*
Popular Science Award	*PS Awards LANL Web*	*O/N 98*	*12*
Records Management	*CIC-10: "Knowledge Is Our Most Important Product"*	*O/N 98*	*3*
SciSearch	*Research Library Adds Social SciSearch® Database*	*F/M 99*	*5*
Supercomputing	*Blue Mountain Is World's Fastest Computer*	*Dec. 98*	*1*
Virtual Private Network	*Virtual Private Network*	*J/J 99*	*15*
Wildfires	*Forecasting Wildfires and other Crises*	*Apr. 98*	*3*
Windows NT	*Preserving Bandwith on Your Windows NT Network*	*F/M 99*	*8*
World Wide Web	*It's Time for PNC: A Graphics Format You Can Pronounce*	*O/N 98*	*7*
(WWW or Web)	*Lab Web Publishers Invited to IntraLab99*	*F/M 99*	*7*
	Extensible HTML	*J/J 99*	*12*
Year 2000 (Y2K)	*The Year 200 Bug is Hiding*	*F/M 99*	*1*

Customer Support Center . . . . (505) 665-4444, ext. 851, or cichelp@lanl.gov

Because of a wide variety of CIC computing services, numerous facilities are available to address your questions. If you are uncertain whom to call, you can always call the Customer Support Center (CSC). CSC consultants are trained to either answer your question or locate someone who can. To reach the appropriate consultant, dial 665-4444 and make your selection from the following choices:

Option 1: New user topics including e-mail, passwords, registration, and World Wide Web
Option 2: Enterprise Information Applications such as Travel, Time and Effort, and Purchase cards
Option 3: Scientific computing, storage systems, and networking
Option 4: Classroom instruction and training
Option 5: Desktop Consulting for PC and Macintosh software and network configurations.

Consulting Via E-Mail

Customer Support ..... cichelp@lanl.gov
Scientific and engineering computing ..... consult@lanl.gov
Administrative and business computing .... eiaconsult@lanl.gov
Passwords and registration ..... validate@lanl.gov
Macintosh computing ..... Mac-help@lanl.gov
PC computing ..... PC-help@lanl.gov
UNIX computing ..... UNIX-help@lanl.gov

Other Useful Numbers

Advanced Computing Laboratory 665-4530
Central Computing Facility 667-4584
Network Operations Center ..... noc@lanl.gov or 667-7423
Telephone Services Center 667-3400

Los Alamos
NATIONAL LABORATORY

Since 1990, the Distributed Computing (CIC-8) High Performance Graphics and Scientific Visualization Team has worked to build a facility where visualization hardware and software are the main focus. The visualization theater shown in this photo is a collaboration between the Applied Theoretical and Computational Physics (X) and Computing, Information, and Communications (CIC) divisions. The theater, or "collabatory," is part of the Numerical Environment for Weapons Simulation, or NEWS, program. For more information, contact Robert Gurulé , at 505-665-3538 or e-mail rgurule@lanl.gov

Los Alamos
NATIONAL LABORATORY

BITS is published monthly to highlight recent computing and communications activities within the Laboratory. We welcome your suggestions and contributions.

BITS may be accessed electronically at this URL:
http://www.lanl.gov/orgs/cic/computingatlanl/

LALP-99-42 (8/99)

Nonprofit organization
US Postage
PAID
Los Alamos, NM
Permit No. 107